Extractors for Images of Varieties

Zeyu Guo[*]
CSE Department
Ohio State University
zguotcs@gmail.com

Ben Lee Volk
Efi Arazi School of Computer Science
Reichman University
benleevolk@gmail.com

Akhil Jalan
Computer Science Department
University of Texas at Austin
akhiljalan@utexas.edu

David Zuckerman[†]
Computer Science Department
University of Texas at Austin
diz@utexas.edu

Abstract

We construct explicit deterministic extractors for *polynomial images of varieties*, that is, distributions sampled by applying a low-degree polynomial map $f : \mathbb{F}_q^r \to \mathbb{F}_q^n$ to an element sampled uniformly at random from a k-dimensional variety $V \subseteq \mathbb{F}_q^r$. This class of sources generalizes both *polynomial sources*, studied by Dvir, Gabizon and Wigderson (FOCS 2007, Comput. Complex. 2009), and *variety sources*, studied by Dvir (CCC 2009, Comput. Complex. 2012).

Assuming certain natural non-degeneracy conditions on the map f and the variety V, which in particular ensure that the source has enough min-entropy, we extract almost all the min-entropy of the distribution. Unlike the Dvir–Gabizon–Wigderson and Dvir results, our construction works over large enough finite fields of arbitrary characteristic. One key part of our construction is an improved deterministic rank extractor for varieties. As a by-product, we obtain explicit Noether normalization lemmas for affine varieties and affine algebras.

Additionally, we generalize a construction of affine extractors with exponentially small error due to Bourgain, Dvir and Leeman (Comput. Complex. 2016) by extending it to all finite prime fields of quasipolynomial size.

[*] Supported in part by a Simons Investigator Award (#409864, David Zuckerman).

[†] Supported in part by NSF Grants CCF-1705028 and CCF-2008076, a Simons Investigator Award (#409864), and the Center of Mathematical Sciences and Applications at Harvard University.

Contents

1 Introduction

Randomness is a powerful resource in computing. There are many useful randomized algorithms, and randomness is provably necessary in cryptography and distributed computing. Naturally, these uses of randomness assume access to uniformly random bits. However, it can be expensive or impossible to obtain such high-quality randomness. A randomness extractor converts low-quality randomness into high-quality randomness.

Low-quality random sources can arise in several ways. First, natural sources of randomness may be defective. Second, in cryptography, if an adversary gains information about a string, then conditioned on this information, the string is weakly random. Third, in constructing pseudorandom generators, a similar situation arises when we condition on the state of the computation. Besides the computer science motivation, randomness extraction questions are natural mathematically.

We model a weak source as a class $\mathcal{D}$ of distributions over a finite set Ω. A randomness extractor for $\mathcal{D}$ is a deterministic function that extracts randomness from any distribution in $\mathcal{D}$.

Definition 1.1. *An extractor for a class $\mathcal{D}$ of distributions with error ε, or an ε-extractor, is a function $\mathsf{Ext} : \Omega \to B$ such that for any $D \in \mathcal{D}$, the distribution $\mathsf{Ext}(D)$ is ε-close, in statistical distance, to the uniform distribution over B.*

Typically the codomain B will be $\{0,1\}^m$.

The most general class of distributions is the set of distributions with high min-entropy, i.e., distributions that do not place much probability on any string. However, it is not hard to show that it is impossible to extract from such sources. It is possible to extract using an auxiliary seed, and there are many applications of such seeded extractors (see [Vad12] for a survey). It is also possible to extract from two independent general weak sources (e.g., [CZ19]). However, if we want to avoid adding a seed and only have one source, we must restrict the class of distributions further.

Various models of weak sources have been studied. It is not hard to show that if there are not too many distributions in the class, then most functions are extractors with excellent parameters. Of course, we really want efficiently-computable extractors.

Models of weak sources tend to be either complexity-theoretic or algebraic. In this work, we focus on *algebraic sources.* That is, we consider distributions over subsets Ω which have a "nice" algebraic structure.

1.1 Algebraic Sources of Randomness

Suppose $\mathbb{F}$ is a finite field and $\Omega = \mathbb{F}^n$. The simplest class of algebraic sources is the set of *affine sources.* An affine source is simply the uniform distribution over an affine subspace $V \subseteq \mathbb{F}^n$ of dimension k. Note that since $|V| = |\mathbb{F}|^k$, the single parameter k also determines the min-entropy of the uniform distribution over the source.

Gabizon and Raz [GR08] constructed an explicit extractor $\mathsf{Ext} : \mathbb{F}^n \to \mathbb{F}^{k-1}$, assuming the field size is bounded from below by a large enough polynomial in n. For a large enough field size q, their construction extracts almost all of the randomness from the source and has error $\varepsilon = 1/\mathrm{poly}(q)$.

The last feature is slightly undesirable, as ideally, one would like the error to decrease exponentially with k, the dimension of the source. Such a construction was given by Bourgain, Dvir and Leeman [BDL16], albeit their construction requires the field size to be slightly super-polynomial in n, and only works for certain fields.

Over smaller fields, constructing affine extractors for small min-entropy is a more challenging task. Further, it is possible to show that any function $f : \mathbb{F}_2^n \to \mathbb{F}_2$ is constant on some affine subspace of dimension $\Omega(\log n)$ (see, e.g., Lemma 6.7 of the arXiv version of [ABGS21]), and thus one cannot hope to extract even a single bit when the min-entropy is smaller than $\log n$ (compare this with the fact that over large fields, the Gabizon–Raz extractor works for any k).

Bourgain [Bou07] constructed an extractor that works over $\mathbb{F}_2$ for min-entropy $k = cn$ for a small constant c. This result was slightly improved by Yehudayoff [Yeh11] and Li [Li11]. Li [Li16] then presented a much improved construction which works when the min-entropy is as small as $k = \log^C(n)$ for some constant C, which was improved by [CGL21] to $k = \log^{1+o(1)}(n)$. However, one drawback of the last two constructions is that the error parameter ε is either constant or polynomially small, whereas one would hope for it to be exponentially small in k, as in the earlier constructions of Bourgain, Yehudayoff and Li.

There are several natural ways to generalize affine sources, but some care is needed when defining those generalizations. As we remarked earlier, for an affine subspace, the single parameter k determines its size and hence the min-entropy of the corresponding source. For more complicated algebraic sets, however, as we shall now see, there are multiple parameters controlling their "complexity," and the connection between those parameters and the min-entropy of the source is not always obvious.

Dvir, Gabizon and Wigderson [DGW09] considered *polynomial sources*, which are defined by applying a low-degree polynomial map $P : \mathbb{F}^k \to \mathbb{F}^n$ on a uniformly random input from $\mathbb{F}^k$. (Note that affine sources are a special case of polynomial sources when the degree equals one.) They further impose the algebraic condition that the Jacobian matrix of the map is of full rank, which in particular guarantees that the min-entropy of the source is high, assuming the characteristic of the field is large enough. The field size required by the construction of [DGW09] is $\mathrm{poly}(k, d, n)^k$.

Dvir [Dvi12] studied a different generalization called *variety sources*, which are uniform distributions over sets $V \subseteq \mathbb{F}^n$ that are the common zeros of a set of low-degree polynomials. Varieties also have an associated concept of dimension, but unlike the affine case, over finite fields having a large dimension does not guarantee by itself that the set V is large, and thus this condition must be imposed explicitly. Dvir presented two constructions. The first requires exponentially large fields and works for any dimension k. The second requires the variety to have size larger than $|\mathbb{F}|^{n/2}$, but the field size depends only polynomially on the degree d of the polynomials defining V.

Over $\mathbb{F}_2$, the situation is much more mysterious. This setting is well motivated, since it turns out that explicit constructions of extractors (or even dispersers) for varieties with various parameters would imply new circuit lower bounds. Golovnev, Kulikov and Williams [GKW21] proved multiple such results. One is that explicit extractors for varieties of size at least $2^{\varepsilon n}$ defined by constant degree polynomials would imply lower bounds for general circuits of the form Cn for larger constants C than what is currently known. They also showed that extractors for varieties of size at least $2^{0.99n}$ defined by polynomials of degree at most $n^{0.01}$ would imply super-linear lower bounds for boolean circuits of depth $O(\log n)$, a long-standing challenge in complexity theory (see also [HR15]).

As for constructions over $\mathbb{F}_2$, Li and Zuckerman [LZ19] showed how to use correlation bounds against low-degree polynomials to obtain extractors for variety sources defined by degree d polynomials for $d = O(1)$ and size at least $2^{(1-c_d)n}$ for some constant c_d that depends on d. Remscrim [Rem16] proved that the majority function is an extractor for varieties defined by polynomials of degree at most n^α and size at least 2^{n-n^β}, assuming $\alpha + \beta < 1/2$. Thus, all the known constructions are not strong enough to imply new circuit lower bounds.

1.2 Our Results

1.2.1 Extractor for Polynomial Images of Varieties

In this paper, we study the class of *polynomial images of varieties*, which generalizes both variety sources and polynomial sources. Informally, the source is specified by a variety $V \subseteq \mathbb{F}^r$ and a polynomial map $f : V \to \mathbb{F}^n$, and a sample from the source is a random variable X computed by uniformly at random picking an element $x \in V$ and outputting $f(x)$. We would like to construct an efficient extractor $\mathsf{Ext} : \mathbb{F}^n \to \{0,1\}^m$ that has small error ε and large output length m. The largest m we can hope for is the min-entropy of the input, which is approximately $k \log q$, where $q = |\mathbb{F}|$ and k is the dimension of the variety V (see Section 4 for a definition of this notion). Our main result is a construction of an extractor with $m \approx k \log q$.

Formally defining such sources takes some care, since varieties and their associated complexity parameters are easier to define over algebraically closed fields. As in previous work, we further need to assume some natural non-degeneracy conditions on the variety V and the map f. We now describe those sources in more detail.

Polynomial images of variety sources. Let $\mathbb{F}$ be a field. For $h_1, \ldots, h_s \in \mathbb{F}[X_1, \ldots, X_n]$, define

$$\mathcal{L}_{h_1,\ldots,h_s,\mathbb{F}} := \{c_0 + c_1 h_1 + \cdots + c_s h_s : c_0, \ldots, c_s \in \mathbb{F}\} \subseteq \mathbb{F}[X_1, \ldots, X_n],$$

i.e., $\mathcal{L}_{h_1,\ldots,h_s,\mathbb{F}}$ is the linear span of $h_1, \ldots, h_s$ and 1 over $\mathbb{F}$.

Denote by $\overline{\mathbb{F}}$ the algebraic closure of $\mathbb{F}$. An *affine variety* $V \subseteq \overline{\mathbb{F}}^n$ over $\mathbb{F}$ is the set of common zeros of a set of polynomials in $\mathbb{F}[X_1, \ldots, X_n]$. Two parameters naturally associated with a variety V are its *dimension*, denoted $\dim V$, which equals the length of the maximal chain with respect to inclusion of distinct irreducible subvarieties, and its *degree*, denoted $\deg V$, which is the number of intersection points of the variety with an affine subspace of codimension $\dim V$ in general position (we refer to Section 4 for more formal definitions).

Definition 1.2 ((n, k, d) algebraic source)**.** *Let $n, d \in \mathbb{N}^+$ and $k \in \mathbb{N}$. We say a distribution D over $\mathbb{F}_q^n$ is an (n, k, d)* algebraic source over $\mathbb{F}_q$ *if there exist $r \in \mathbb{N}$, an affine variety $V \subseteq \overline{\mathbb{F}}_q^r$ over $\mathbb{F}_q$, polynomials $h_1, \ldots, h_s \in \mathbb{F}_q[X_1, \ldots, X_r]$ with $\deg h_1 \geq \cdots \geq \deg h_s$, and $f_1, \ldots, f_n \in \mathcal{L}_{h_1,\ldots,h_s,\mathbb{F}_q}$ such that $D = f(U_{V(\mathbb{F}_q)})$, where $f : \overline{\mathbb{F}}_q^r \to \overline{\mathbb{F}}_q^n$ is the polynomial map defined by $f_1, \ldots, f_n$, and $U_{V(\mathbb{F}_q)}$ is the uniform distribution over $V(\mathbb{F}_q) := V \cap \mathbb{F}_q^r$, and further, the following conditions hold:*

1. *At least one irreducible component of V of dimension $\dim V$ is absolutely irreducible.*

2. *For every irreducible component V_0 of dimension $\dim V$ that is absolutely irreducible, the dimension of $\overline{f(V_0)}$ is at least k, where $\overline{f(V_0)} \subseteq \overline{\mathbb{F}}_q^n$ denotes the closure of $f(V_0)$, i.e., the smallest affine variety over $\mathbb{F}_q$ containing $f(V_0)$.*

3. *$\deg V \cdot \prod_{i=1}^k \deg h_i \leq d$.*[1]

In addition, we say D is an irreducible *(n, k, d) algebraic source over $\mathbb{F}_q$ if V can be chosen to be irreducible. We say D is a* minimal *(n, k, d) algebraic source over $\mathbb{F}_q$ if V can be chosen to have dimension k. Finally, we say D is an* irreducibly minimal *(n, k, d) algebraic source over $\mathbb{F}_q$ if V can be chosen to be irreducible of dimension k.*

[1] Note that $\dim \overline{f(V)} \geq k$ by previous conditions. So we necessarily have $s \geq k$ and $\deg h_i \geq 1$ for $i \in [k]$. This also implies $\deg V \leq d$.

The conditions in Definition 1.2 may look a bit contrived at first glance. However, as we now explain, they are quite natural, and indeed some form of them, as observed in previous work, is necessary.

The third condition is simply a convenient way to "pack" multiple "complexity" parameters of the components of the source that arise in the analysis. That is, d is a single complexity parameter that, in particular, bounds the degree of the variety V and the product of degrees of the polynomial map f. Having d as a single parameter simplifies the statements of our theorems and clarifies the dependence between the various parameters: the larger d is, the larger the field size we require and the smaller the output length of the extractor.

The purpose of the first two conditions is to guarantee that our source has enough min-entropy. As observed in previous work [DGW09, Dvi12], it is quite easy to come up with simple varieties V (even of high dimension) or polynomial maps f (even of low degree) such that sources arising as $f(V)$ would have very few points in $\mathbb{F}_q^n$, so that there will be little to no randomness to extract.

The first condition is analogous to (and, as shown in Appendix C, roughly equivalent to) Dvir's [Dvi12] condition that the variety V contains enough points in $\mathbb{F}_p^n$. The second condition is analogous to (and, over fields of large characteristic, implied by) the full-rank Jacobian condition of Dvir, Gabizon and Wigderson [DGW09]. Thus, not only is some form of conditions 1 and 2 necessary for proving any meaningful results, but moreover, these conditions naturally generalize the conditions imposed by previous related works.

Finally, we note that the name "(n, k, d) algebraic sources" suppresses the dependence on the parameter r in the definition, which is the ambient dimension in which the variety V lies. This is because our result, stated next, has no dependence on r. Even in the case where r is very large with respect to n, k and d, our results only depend on the latter three parameters. Further, note that when r is very large, $\dim V$ can also be very large compared with n and k. However, as the definition hints, we will reduce this case to the case where $\dim V = k$.

We can now state our main theorem.

Theorem 1. *Let $n, d \in \mathbb{N}^+$, $k \in \mathbb{N}$, and $\varepsilon \in (0, 1/2]$. Let q be a power of a prime p. Suppose $q \geq (nd/\varepsilon)^c$, where $c > 0$ is a large enough absolute constant. Then there exists an explicit ε-extractor $\mathsf{Ext} : \mathbb{F}_q^n \to \{0,1\}^m$ for (n, k, d) algebraic sources over $\mathbb{F}_q$ with output length $m \geq k \log q - 4 \log\log p - O(\log(nd/\varepsilon))$.*

It can be shown that any (n, k, d) algebraic source D over $\mathbb{F}_q$, where $q \geq (kd)^c$ for a sufficiently large constant $c > 0$, is (close to) a distribution with min-entropy at least $k \log q - O(\log d)$. Moreover, this estimate of the min-entropy is tight up to an additive term $O(\log d)$ if D is not an $(n, k+1, d)$ algebraic source over $\mathbb{F}_q$. See Lemma 7.10 and Proposition 7.11. Therefore, the extractor in Theorem 1 extracts most of the min-entropy from (n, k, d) algebraic sources. In addition, Theorem 1 works over finite fields of any characteristic, while the extractors by Dvir, Gabizon, and Wigderson [DGW09] and Dvir [Dvi12] require large enough characteristics.

As is standard in the literature, by "explicit" we mean that the output of the extractor is computable in time $\mathrm{poly}(n, \log q)$ (note that the input length to the extractor is $n \log q$).

Along the way to proving Theorem 1, we construct several other algebraic pseudorandom objects which are interesting on their own. We mention some of these constructions when we give an overview of our construction in Section 1.3.

1.2.2 Affine Extractors for Quasipolynomally Large Fields with Exponentially Small Error

Recall that an explicit affine extractor is an efficiently computable function $\mathsf{Ext} : \mathbb{F}^n \to \mathbb{F}^m$ such that for every affine subspace $V \subseteq \mathbb{F}^n$ of dimension k, and a random variable X uniformly sampled from V, $\mathsf{Ext}(X)$ is close to the uniform distribution over $\mathbb{F}^m$. We would like m to be as close to k as possible and, ideally, the error parameter ε to be exponentially small in k.

As mentioned earlier, the extractor of Gabizon and Raz [GR08] achieves $m = k - 1$ and error ε only polynomially small in the field size q. In particular, the error does not decrease with k. Bourgain, Dvir and Leeman [BDL16] constructed an affine extractor with m arbitrarily close to $k/2$ and error $q^{-\Omega(k)}$. However, their construction requires q to be slightly super-polynomial in n, namely $q = n^{\Omega(\log \log n)}$, and furthermore only works for "most" prime fields $\mathbb{F}_q$. We improve the analysis of their construction and present a construction with identical parameters that works for *all* prime fields, assuming $q = n^{\Omega(\log \log n)}$.

Theorem 2. *For every $0 < \beta < 1/2$, there exists a constant C such that the following holds: Let $k \leq n$ be integers and $\mathbb{F}$ be a prime field of size $q \geq n^{C \log \log n}$. Let $m = \beta k$. There exists an efficiently computable function $E : \mathbb{F}^n \to \mathbb{F}^m$ which is an affine extractor for min-entropy k with error $q^{-\Omega(k)}$.*

1.3 Techniques

Our construction from Theorem 1 combines several techniques used in previous related constructions, as well as several new ideas which are required to successfully apply these techniques. It is convenient to think of the construction as proceeding in several steps.

Preliminary step: decomposing the sources. Our definition for algebraic sources (Definition 1.2) is quite general, and it is convenient to work with slightly "nicer" sources. We start by approximating general (n, k, d) algebraic sources as convex combinations of *irreducibly minimal* (n, k, d) algebraic sources. Recall that this means that the variety V is irreducible and has dimension k.

This step is done in Section 7: we first decompose a general source into a convex combination of irreducible sources in a manner that follows naturally from the decomposition of V itself as a union of irreducible components. We then decompose an irreducible source into irreducibly minimal sources roughly by intersecting it with a linear space of the appropriate dimension. Both parts of the arguments incur a small error.

First step: extracting a short seed. Having reduced to the case of irreducibly minimal sources, we first design an extractor that extracts a small number of bits from the source. One commonly used technique for doing that is to show that the source is an ε-biased distribution, i.e., a distribution whose nontrivial Fourier coefficients are all small. Similar methods work when the source is close to such a distribution or to a convex combination of such distributions. Analyzing and bounding the Fourier coefficients is often done using bounds on exponential sums from algebraic geometry, such as Bombieri's estimate (Theorem 4.7). We follow this general paradigm as well.

However, the case where the field characteristic is small presents some unique challenges to overcome. We first prove an extension of Bombieri's theorem for small characteristic p. This extension bounds the corresponding exponential sums save for possibly a small set of "bad" characters.

Hence, we then define and study a more general class than ε-biased distributions: (ε, e)-biased distributions, which are distributions in which all but at most e of the Fourier coefficients have absolute value at most ε. We show that the sources we consider are close to convex combinations of such distributions (for meaningful values of ε and e), and construct extractors for such distributions.

Previously, the XOR lemma has been used to construct extractors for ε-biased sources; see, e.g., Rao [Rao07]. We extend these ideas to the more general and challenging setting of (ε, e)-biased distributions. On the technical level, we construct explicit functions $f : \mathbb{F}_p^n \to \mathbb{F}_p^t$ with the following properties: for every nontrivial character ψ of $\mathbb{F}_p^t$, both the L_1 and the L_∞ norms of the Fourier transform of $\psi \circ f$ (which is a function from $\mathbb{F}_p^n$ to $\mathbb{C}$) are upper bounded by sufficiently small quantities. We in fact present two constructions of such functions f. The first is based on standard error-correcting codes over $\mathbb{F}_p$, and the second is an improved construction based on *rank-metric codes*. Those constructions appear in Section 3.2.

Second step: applying a seeded extractor. Having extracted a small number of bits, we wish to use them as a *seed* in an application of a seeded extractor on the source to extract almost all the min-entropy. The challenge, of course, is that the seed is correlated with the source, whereas a seeded extractor requires the seed to be independent of the source. Techniques for dealing with these problems were developed in [GRS06, GR08], as this is also the general methodology in their extractor constructions. This is done by analyzing the conditional distribution of the source conditioned on any possible output of the seeded extractor with a fixed seed, and showing that it maintains some nice properties. We first analyze the case where the image $f(V)$ of the polynomial map is of full rank inside $\mathbb{F}^k$, using the *effective fiber dimension theorem*. We then consider the general case. In order to reduce to that case, we apply a *rank extractor* for varieties, a notion we define and develop in this work, building upon previous work which developed rank extractors for linear spaces.

Rank extractor for varieties. Let $V \subseteq \mathbb{F}^n$ be a k-dimensional variety. We would like to obtain a map $E : \mathbb{F}^n \to \mathbb{F}^k$ which "extracts" all the rank from V, in the sense that $E(V) \subseteq \mathbb{F}^k$ is k-dimensional. The first obvious challenge is that $E(V)$ need not necessarily be a variety. It is thus natural in this case to consider the closure of $E(V)$ in $\overline{\mathbb{F}}^n$ where $\overline{\mathbb{F}}$ is the algebraic closure of $\mathbb{F}$.

Previous work has considered the case where V is a linear subspace. In this case, observe that if E is linear, then $E(V)$ is also a linear subspace. However, there clearly cannot be a single map E that preserves the dimension of all linear subspaces, as given any fixed E, one could take V to be the kernel of E. Therefore, a natural relaxation is to consider *seeded* linear rank extractors, which are collections of linear maps $E_1, \ldots, E_t$ such that for every V, most of the maps preserve the dimension. Such objects were first defined and constructed by Gabizon and Raz [GR08]. Improved and optimal parameters (in terms of the "seed length," i.e., the number of maps) were obtained by Forbes and Shpilka [FS12], and a systematic study of these objects appears in [FG15].

In this work, we observe that seeded linear rank extractors for extractors are also seeded linear rank extractors for varieties (see Section 5). The key insight is that rank extractors (for linear subspaces) preserve the dimensions of the tangent spaces at nonsingular points of the variety, which turns out to be a sufficient criterion.

Linear rank extractors are very useful because they enable us to condense sources that are not full-rank to full-rank sources without increasing the degrees of the polynomial maps. However, it turns out that it is also possible to construct *deterministic* rank extractors for varieties, which we do in Section 6. Such extractors are obviously not linear maps, although in our constructions, they are polynomials of fairly small degrees (polynomial in n and in the degree d of the variety). We remark that Dvir [Dvi12] constructed such an extractor for one-dimensional varieties, and his extractor is a polynomial of degree exponential in n. In addition, Dvir, Gabizon and Wigderson [DGW09] constructed rank extractors for polynomial sources using a different technique.

Our construction adapts the construction of Dvir, Kollár and Lovett [DKL14], who constructed different pseudorandom objects called *variety evasive sets*. By modifying their proof, we are able to show that a similar construction yields a deterministic rank extractor for varieties. This essentially follows because their map φ satisfies the property that for every low-degree variety V and every point $b \in \mathbb{F}^k$, the intersection $\varphi^{-1}(b) \cap V$ is a finite set. Dvir, Kollár and Lovett prove it only for the case $b = \mathbf{0}$, but it is not hard to extend it to general b.

Explicit Noether normalization lemmas. As a by-product of the above construction of deterministic rank extractors for varieties, we prove explicit *Noether normalization lemmas* for affine varieties and affine algebras. The Noether normalization lemma [Noe26, Nag62] is a classical result in commutative algebra and algebraic geometry, which states that any affine variety of dimension k admits a surjective *finite morphism* to an affine space of dimension k. We show that the construction in [DKL14] in fact gives a direct construction of such a finite morphism. In contrast, the textbook proof of Nagata [Nag62] is iterative and uses polynomials of degrees that are at least doubly exponential in the number of steps of the iteration.

Our proof is inspired by a geometric argument of Kollár, Rónyai and Szabó [KRS96]. See Section 11 and Appendix D for more details.

Affine extractors with exponentially small error. Our proof of Theorem 2 follows a very similar route to the proof of the main theorem of Bourgain, Dvir and Leeman [BDL16], who constructed such an extractor for prime fields $\mathbb{F}_q$ for "typical" primes q. Our main contribution is an improved number-theoretic lemma (Proposition 10.3) which shows how to find n distinct integers $d_1, \ldots, d_n$ with desirable number theoretic properties. The proof then proceeds by estimating the Fourier coefficient of the distribution obtained by applying our extractor to a linear subspace using an exponential sum estimate of Deligne, much in the same way as [BDL16].

1.4 Comparison with Previous Work

The two works closest to ours are by Dvir [Dvi12] and Dvir, Gabizon and Wigderson [DGW09], both of which construct extractors for sources with algebraic structures.

As mentioned earlier, Dvir, Gabizon and Wigderson [DGW09] study *polynomial sources*, defined by picking an element $x \in \mathbb{F}_q^k$ uniformly at random and applying a polynomial map $f : \mathbb{F}_q^k \to \mathbb{F}_q^n$ of degree at most d. This is a special case of the sources we consider when the variety V is taken to be $\mathbb{F}_q^k$.

They further add the non-degeneracy condition that the *Jacobian* of the mapping f, namely, its matrix of partial derivatives, has full rank. This in particular guarantees that the source has a high enough min-entropy. Their main theorem gives an explicit extractor that outputs a constant fraction of the min-entropy over prime fields $\mathbb{F}_p$ of cardinality $\mathrm{poly}(n, d)^{Ck}$ for some constant C. Our construction in Theorem 1, on the other hand, works for a larger class of sources, outputs almost all the min-entropy, and works over finite fields of small characteristics as well.

Dvir [Dvi12] considers *variety sources*, which he defines as uniform distributions over sets of the type $\{x : f_1(x) = f_2(x) = \cdots = f_t(x) = 0\}$ in $\mathbb{F}_p^n$, where $\deg f_i \leq d$ for all i. These sources are also a special case of the type of sources we consider. One should note, however, the different usage of the term "degree" in our definitions: Dvir always refers to the degree $\deg f_i$ of the polynomials which define the variety V, whereas we refer to the degree $\deg V$ of V as an affine variety, which is often much larger.

Assuming $\dim V = k$ and $|V| \geq p^{k-c}$ for some small constant $c > 0$, Dvir's extractor [Dvi12] outputs a constant fraction of the min-entropy over prime fields of characteristic $p > d^{Cn^2}$ for some constant C. Again, Dvir uses the parameter d differently than we do in Theorem 1. In particular, in our construction, the field size q is only polynomial in the parameter d (but d might be exponential in n).

As mentioned in the discussion after Definition 1.2, our assumptions are weaker than those of [DGW09] and [Dvi12]. Thus, as our sources is more general, the characteristic in our results can be arbitrary, and our conclusions are stronger (since we extract more output bits), it follows that in particular our result subsumes the extractors of [DGW09] and [Dvi12].

Dvir [Dvi12] also presents a different construction that outputs a very small number of bits from very large varieties over small fields. This construction is incomparable with our results.

On the more technical level, we discuss a particular feature of our proof that distinguishes it from [DGW09, Dvi12] and, in particular, allows us to extend the output length.

For simplicity, consider the case of $(1, 1, d)$ algebraic sources. As mentioned in Section 1.3, we first prove an extension of Bombieri's estimate that holds even if the characteristic p is small: if p is small, this result implies that a $(1, 1, d)$ algebraic source D over $\mathbb{F}_q$ is a convex combination of (ε, d)-biased sources. That is, we allow a few large Fourier coefficients. Then we use the machinery developed in Section 3 to extract randomness from D. On the other hand, if p is large enough, then D has no large nontrivial Fourier coefficients; it is ε-biased. In this case, the XOR lemma is sufficient, as argued in [DGW09, Dvi12].

To apply Bombieri's estimate to a high-dimensional affine variety V, we follow [DGW09, Dvi12] and decompose V into a family of affine curves C_i such that the polynomial f that does not vanish identically on V still does not vanish on most C_i.

In [DGW09], this is achieved using an argument based on the Jacobian criterion for algebraic independence, but it works only when the characteristic p is large. Instead of using this argument, we use the decomposition of (n, k, d) algebraic sources into irreducibly minimal (n, k, d) algebraic sources proved in Section 7, whose proof is based on the effective fiber dimension theorem (Theorem 4.10) and works for any characteristic.

The last idea we introduce is the use of the effective Lang–Weil bound (Theorem 4.6), which allows us to extract almost $\log q$ bits. To explain the idea, consider an affine variety $V \subseteq \mathbb{A}^n_{\mathbb{F}_q}$ and write $V(\mathbb{F}_q)$ as a disjoint union of $C_i(\mathbb{F}_q)$ for a family of affine curves C_i over $\mathbb{F}_q$. Let f be a low-degree polynomial and assume for simplicity that f is non-constant on every C_i. Let χ be a nontrivial character of $\mathbb{F}_q$. The following win-win argument was used in [DGW09] to bound the bias $\delta := \left|\mathbb{E}_{x \in V(\mathbb{F}_q)}[\chi(f(x))]\right|$: For a curve C_i, if $|C_i(\mathbb{F}_q)|$ is small, say $|C_i(\mathbb{F}_q)| \leq \Delta$ for some threshold Δ, then its contribution to the bias δ is small assuming that V has many rational points. On the other hand, if $C_i(\mathbb{F}_q) > \Delta$, then Bombieri's estimate (Lemma 8.3), together with the fact that

$$\left|\mathop{\mathbb{E}}_{x \in C_i(\mathbb{F}_q)}[\chi(f(x))]\right| = \frac{\left|\sum_{x \in C_i(\mathbb{F}_q)}[\chi(f(x))]\right|}{|C_i(\mathbb{F}_q)|} \leq \frac{\left|\sum_{x \in C_i(\mathbb{F}_q)}[\chi(f(x))]\right|}{\Delta},$$

implies that $\left|\mathbb{E}_{x \in C_i(\mathbb{F}_q)}[\chi(f(x))]\right|$ is small. Considering all curves C_i shows that the bias is small. We note that no information about $|C_i(\mathbb{F}_q)|$ was used in this win-win argument. For this reason, the choice of threshold Δ cannot be too large or too small, and the resulting extractors only extract a constant fraction of $\log q$ bits. To improve the output length, we observe that the effective Lang–Weil bound (Theorem 4.6) together with Lemma 7.2 gives more information about $|C_i(\mathbb{F}_q)|$. In particular, for an irreducible affine curve C, the number $|C(\mathbb{F}_q)|$ is either close to q or very small, depending on whether C is absolutely irreducible. Exploiting this fact yields an explicit construction of deterministic extractors that output almost $\log q$ bits.

1.5 Open Problems

While improving the dependence on any of the parameters in our construction remains an open problem, in our opinion, the main challenge is reducing the field size. In our construction for polynomial images of varieties (Theorem 1), we require field size $\mathrm{poly}(n, 1/\varepsilon, d)$. We stress that for certain varieties, d can be exponential in n (although it is by no means necessarily so). Can we construct extractors for significantly smaller fields, perhaps even constant size?

As mentioned above, over very small fields, such as $\mathbb{F}_2$, certain Ramsey-theoretic lower bounds imply that constructions such as ours that work for any min-entropy cannot exist. A key reason to study $\mathbb{F}_2$ is that explicit extractors with certain parameters imply new circuit lower bounds.

In our construction of new affine extractors (Theorem 2), we obtain a field size that is slightly super-polynomial in n. It is a very appealing open problem to reduce the field size to a polynomial in n.

A related problem is reducing the degree of our deterministic rank extractor. In Section 6, we construct a deterministic rank extractor for varieties whose degree is $\mathrm{poly}(n, d)$ for degree d varieties. Reducing the degree, perhaps to depend only on d, would help lower the field size requirement for the extractor for polynomial images of varieties to depend only on the degree.

We end with two general questions. Can our constructions or techniques help in designing extractors for larger and more general classes of sources, either algebraic or complexity-theoretic? Do our constructions have any complexity-theoretic implications, such as lower bounds for certain models of computation?

2 Notations and Preliminaries

Let $\mathbb{N} = \{0, 1, \dots\}$, $\mathbb{N}^+ = \{1, 2, \dots, \}$, and $[n] = \{1, 2, \dots, n\}$ for $n \in \mathbb{N}$. Write $\mathbb{Z}_n$ for the cyclic group $\{0, 1, \dots, n-1\}$ with addition modulo n.

The cardinality of a set S is denoted by $|S|$. We also use $|c|$ to denote the absolute value of a number $c \in \mathbb{C}$. Denote by $\log x$ the base 2 logarithm of x, and by $\ln x$ the natural logarithm of x. For sets A and B, denote by $A \setminus B$ the set difference $\{x \in A : x \notin B\}$. The restriction of a map $f : A \to B$ to a subset $A' \subseteq A$ is denoted by $f|_{A'}$, which is a map from A' to B.

A *formal Laurent series* over a field $\mathbb{F}$ has the form

$$h(T) = c_i T^i + c_{i+1} T^{i+1} + \dots$$

where $i \in \mathbb{Z}$ and $c_j \in \mathbb{F}$ for $j \geq i$. Denote by $\mathrm{ord}(f)$ the least degree of the terms that appear in f, i.e., $f = c_0 T^{\mathrm{ord}(f)} + c_1 T^{\mathrm{ord}(f)+1} + \cdots$ where $c_0 \neq 0$. If $f = 0$, then let $\mathrm{ord}(f) = +\infty$. The set of formal Laurent series over $\mathbb{F}$ is a field, denoted by $\mathbb{F}((T))$. We say $f \in \mathbb{F}((T))$ is a *formal power* *series* over $\mathbb{F}$ if $\mathrm{ord}(f) \geq 0$. The set of formal power series over $\mathbb{F}$ is a subring of $\mathbb{F}((T))$, denoted by $\mathbb{F}[[T]]$.

We write $x \sim D$ if x is sampled from a distribution D. The *support* of a distribution D over a finite set Ω is $\mathrm{supp}(D) := \{a \in \Omega : \Pr[D = a] \neq 0\}$. For an event A that occurs with a nonzero probability under a distribution D, write $D|_A$ for the distribution of D conditioned on A. The product distribution of two distributions D, D' is denoted by $D \times D'$. The *statistical distance* between two distributions D, D' over a finite set Ω is defined to be

$$\Delta(D, D') := \max_{A \subseteq \Omega} |\Pr[D \in A] - \Pr[D' \in A]|.$$

Two distributions D and D' are *ε-close* if their statistical distance is at most ε, and we write $D =_\varepsilon D'$ for this statement.

The uniform distribution over a finite set S is denoted by U_S. For $n \in \mathbb{N}$, denote by U_n the uniform distribution over $\{0,1\}^n$.

The *min-entropy* of a distribution D over a finite set Ω is

$$H_{\min}(D) := -\log(\max_{a \in \Omega} \Pr[D = a]).$$

We say D is a *k-source* if $H_{\min}(D) \geq k$.

Let Ω and B be finite sets, and let $\mathcal{D}$ be a class of distributions over Ω. A function $\mathsf{Ext} : \Omega \to B$ is said to be a *(deterministic) ε-extractor* for $\mathcal{D}$ if $\mathsf{Ext}(D) =_\varepsilon U_B$ for all $D \in \mathcal{D}$. A function $\mathsf{Ext} : \Omega \times \{0,1\}^\ell \to B$ is said to be a *seeded ε-extractor* for $\mathcal{D}$ if $\mathsf{Ext}(D \times U_\ell) =_\varepsilon U_B$ for all $D \in \mathcal{D}$, where $\ell \in \mathbb{N}$ is called the *seed length* of Ext.

Facts about distributions. The following lemmas are standard and can be found in, e.g., [Sha08, Section 2].

Lemma 2.1. *Let $f : A \to B$ be a map between finite sets. Let D and D' be distributions over A. If $D =_\varepsilon D'$, then $f(D) =_\varepsilon f(D')$.*

Lemma 2.2. *Let $D = (D_1, D_2)$ be a joint distribution over a finite product set $A \times B$. Suppose T is a subset of A such that $\Pr[D_1 \in T] \geq 1 - \varepsilon_1$ and $D_2|_{D_1 = a} =_{\varepsilon_2} D_2'$ for some distribution D_2' over B and all $a \in T \cap \mathrm{supp}(D_1)$. Then $D =_{\varepsilon_1 + \varepsilon_2} D_1 \times D_2'$.*

3 Sources with Low Bias and Their Extractors

In this section, we consider several natural extensions of ε-biased sources which are useful for our extractor constructions. We then show how to extract randomness from such sources.

3.1 (ε, e)-Biased Sources

Let A be a finite abelian group and let $\widehat{A}$ denote the dual group of A, that is, the group of characters over A. A distribution D over A is *ε-biased* if $|\mathbb{E}[\chi(D)]| \leq \varepsilon$ for all nontrivial characters $\chi \in \widehat{A}$. This is a standard definition, introduced in [NN93], which has been immensely useful in the construction of extractors and in the theory of pseudorandomness in general.

We now introduce two natural generalizations. We say D is *(ε, e)-biased* if $|\mathbb{E}[\chi(D)]| \leq \varepsilon$ for all but at most e characters $\chi \in \widehat{A}$. And we say D is *strongly (ε, e)-biased* if the set of $\chi \in \widehat{A}$ satisfying

$|\mathbb{E}[\chi(D)]| > \varepsilon$ is contained in an abelian subgroup of A of size at most e. The usefulness of the latter definition will be clear shortly.

Note that if A is a vector space over a finite field of characteristic p, then an (ε, e)-biased distribution is also strongly (ε, p^e)-biased.

We have the following easy observation, which shows that strongly (ε, e)-biased sources generalize affine sources (of low codimension).

Lemma 3.1. *An affine source of codimension at most k in $\mathbb{F}_p^n$ is strongly $(0, p^k)$-biased.*

The following proposition states that (ε, e)-biased sources are (close to) distributions with high min-entropy.

Proposition 3.2. *Let D be an (ε, e)-biased distribution over A. Then D is ε'-close to a k-source if $\varepsilon' \geq 2^k(\varepsilon^2 + |A|^{-1}e)$. In particular, this holds for $k = \min\{2\log(1/\varepsilon), \log|A| - \log e\} - \log(2/\varepsilon')$.*

Proof. View D as a function $D : A \to \mathbb{R}$ such that $D(x) = \Pr[D = x]$ for $x \in A$. For $\chi \in \widehat{A}$, we have $\widehat{D}(\chi) = \mathbb{E}_{x\in A}[D(x)\overline{\chi(x)}] = |A|^{-1}\left(\sum_{x\in A}\Pr[D = x]\cdot\overline{\chi(x)}\right) = |A|^{-1}\overline{\mathbb{E}[\chi(D)]}$. So $|\widehat{D}(\chi)| \leq |A|^{-1}$ for all $\chi \in \widehat{A}$, and $|\widehat{D}(\chi)| \leq |A|^{-1}\varepsilon$ for all but at most e characters $\chi \in \widehat{A}$. By Parseval's identity,

$$\sum_{x\in A}|D(x)|^2 = |A|\cdot\left(\mathop{\mathbb{E}}_{x\in A}\left[|D(x)|^2\right]\right) = |A|\cdot\left(\sum_{\chi\in\widehat{A}}|\widehat{D}(\chi)|^2\right) \leq \varepsilon^2 + |A|^{-1}e.$$

Let $A' \subseteq A$ be the set of $x \in A$ such that $D(x) > 2^{-k}$. Then

$$\varepsilon^2 + |A|^{-1}e \geq \sum_{x\in A}|D(x)|^2 \geq \sum_{x\in A'}|D(x)|^2 \geq \left(\sum_{x\in A'}D(x)\right)2^{-k}.$$

So the total probability mass contributed by $x \in A'$ is bounded by $\varepsilon' = 2^k(\varepsilon^2 + |A|^{-1}e)$. This implies that D is ε'-close to a k-source. □

Next, suppose that A and B are finite groups. We wish to bound the bias of conditional distributions over A (or B), assuming bounds on the bias of a distribution over $A \times B$. We begin with the following technical calculation.

Lemma 3.3. *Let A and B be finite abelian groups. Identifying $\widehat{A} \times \widehat{B}$ with $\widehat{A \times B}$ so that $(\chi, \theta)(x, y) = \chi(x)\theta(y)$ for $(x, y) \in A \times B$ and $(\chi, \theta) \in \widehat{A} \times \widehat{B}$. Let $D = (D_1, D_2)$ be a joint distribution over $A \times B$. For $x \in \mathrm{supp}(D_1)$ and $\theta \in \widehat{B}$, we have*

$$\mathbb{E}[\theta(D_2|_{D_1=x})] = \sum_{\chi\in\widehat{A}}\Pr[D_1 = x]^{-1}\cdot|A|^{-1}\cdot\overline{\chi(x)}\cdot\mathbb{E}[(\chi,\theta)(D)].$$

Proof. Define $\delta_x : A \to \{0, 1\}$ to be the indicator function such that $\delta_x(z) = 1$ if and only if $z = x$. Then

$$\begin{aligned}\mathbb{E}[\theta(D_2|_{D_1=x})] &= \Pr[D_1 = x]^{-1}\cdot\mathbb{E}[\delta_x(D_1)\theta(D_2)]\\ &= \Pr[D_1 = x]^{-1}\cdot\mathbb{E}\left[\left(\sum_{\chi\in\widehat{A}}\widehat{\delta_x}(\chi)\chi(D_1)\right)\theta(D_2)\right]\end{aligned}$$

$$
\begin{aligned}
&= \sum_{\chi \in \widehat{A}} \Pr[D_1 = x]^{-1} \cdot \widehat{\delta_x}(\chi) \cdot \mathbb{E}[\chi(D_1)\theta(D_2)] \\
&= \sum_{\chi \in \widehat{A}} \Pr[D_1 = x]^{-1} \cdot |A|^{-1} \cdot \overline{\chi(x)} \cdot \mathbb{E}[(\chi, \theta)(D)]
\end{aligned}
$$

where the last equality uses the fact $\widehat{\delta_x}(\chi) = \mathbb{E}_{z \in A}[\delta_x(z)\overline{\chi(z)}] = |A|^{-1} \cdot \overline{\chi(x)}$. □

As a corollary, we bound the bias of the marginal distribution D_2 conditioned on any value of D_1.

Corollary 3.4. *Use the notations in Lemma 3.3. Let $\varepsilon, \varepsilon' > 0$. Assume that every character $\chi \in \widehat{A \times B} \cong \widehat{A} \times \widehat{B}$ satisfying $\mathbb{E}[\chi(D)] > \varepsilon$ is contained in the subgroup $\widehat{A} \times \{1\}$. Then with probability at least $1 - \varepsilon'$ over $x \sim D_1$, the conditional distribution $D_2|_{D_1 = x}$ is $|A|\varepsilon/\varepsilon'$-biased.*

Proof. Let T be the set of $x \in A$ satisfying $\Pr[D_1 = x] \leq |A|^{-1}\varepsilon'$. Then $\Pr[D_1 \in T] = \sum_{x \in T} \Pr[D_1 = x] \leq \varepsilon'$. So it suffices to show that $|\mathbb{E}[\theta(D_2|_{D_1=x})]| \leq |A|\varepsilon/\varepsilon'$ holds for every $x \in A \setminus T$ and every nontrivial character θ of B.

Consider $x \in A \setminus T$ and a nontrivial character θ of B. As $x \notin T$, we have $\Pr[D_1 = x] > |A|^{-1}\varepsilon'$. By Lemma 3.3,

$$
\begin{aligned}
|\mathbb{E}[\theta(D_2|_{D_1=x})]| &= \left| \sum_{\chi \in \widehat{A}} \Pr[D_1 = x]^{-1} \cdot |A|^{-1} \cdot \overline{\chi(x)} \cdot \mathbb{E}[(\chi, \theta)(D)] \right| \\
&\leq \left| \sum_{\chi \in \widehat{A}} \overline{\chi(x)} \cdot \mathbb{E}[(\chi, \theta)(D)] \right| / \varepsilon' \\
&\leq \sum_{\chi \in \widehat{A}} |\mathbb{E}[(\chi, \theta)(D)]| / \varepsilon'.
\end{aligned}
$$

Note that for $\chi \in \widehat{A}$, (χ, θ) is not in the subgroup $\widehat{A} \times \{1\}$ since $\theta \in \widehat{B}$ is nontrivial. So $|\mathbb{E}[(\chi, \theta)(D)]| \leq \varepsilon$ by assumption. It follows that $|\mathbb{E}[\theta(D_2|_{D_1=x})]| \leq |A|\varepsilon/\varepsilon'$, as desired. □

3.2 Extraction via the XOR Lemma and Rank-Metric Codes

We need the following form of Vazirani's XOR lemma, taken from [Rao07].

Lemma 3.5 (XOR lemma)**.** *Every ε-biased distribution over a finite abelian group A is $\varepsilon|A|^{1/2}$-close to the uniform distribution over A.*

Large characteristic. Over fields of large characteristic, we use the mod-M function in our constructions as an extractor for low-bias distributions, in a similar manner to [DGW09] and [Dvi12]. We follow the treatment in [Rao07].

For a finite abelian group A and a function $h : A \to \mathbb{C}$, define $\|h\|_1 = \sum_{x \in A} |h(x)|$ and $\|h\|_\infty = \max_{x \in A} |h(x)|$.

Lemma 3.6 ([Rao07, Lemma 4.3])**.** *Let A and B be finite abelian groups. Let D be an ε-biased distribution over A. Suppose $f : A \to B$ is a map such that for every character ψ of B, we have that $\left\|\widehat{\psi \circ f}\right\|_1 \leq \tau$. Then $f(D)$ is ε'-close to $f(U_A)$, where $\varepsilon' = \tau\varepsilon|B|^{1/2}$.*

Lemma 3.7 ([Rao07, Lemma 4.4]). *Let $f : \mathbb{Z}_N \to \mathbb{Z}_M$ be the map sending $a \bmod N$ to $a \bmod M$ for $a \in \{0, 1, \dots, N-1\}$. Let ψ be a character of $\mathbb{Z}_M$. Then $\left\|\widehat{\psi \circ f}\right\|_1 \leq c \log N$, where c is an absolute constant.*

When p is large but $\mathbb{F}_q$ is possibly non-prime, we simply apply the mod-M function to the last $\mathbb{F}_p$-coordinate of $\mathbb{F}_q$ and use the following corollary of Lemma 3.7.

Corollary 3.8. *Let $f : \mathbb{Z}_N^t \to \mathbb{Z}_N^{t-1} \times \mathbb{Z}_M$ be the map that sends $(a_1, \dots, a_{t-1}, a \bmod N)$ to $(a_1, \dots, a_{t-1}, a \bmod M)$ for $(a_1, \dots, a_{t-1}, a) \in \mathbb{Z}_N^{t-1} \times \{0, 1, \dots, N-1\}$. Let ψ be a character of $\mathbb{Z}_N^{t-1} \times \mathbb{Z}_M$. Then $\left\|\widehat{\psi \circ f}\right\|_1 \leq c \log N$, where c is an absolute constant.*

Combining Lemma 3.6 and Corollary 3.8 gives the following lemma, which allows us to extract randomness from ε-biased sources over $\mathbb{F}_q$.

Lemma 3.9. *Let $f : \mathbb{Z}_N^t \to \mathbb{Z}_N^{t-1} \times \mathbb{Z}_M$ be the map in Corollary 3.8. Then for every ε-biased distribution D over $\mathbb{Z}_N^t$, $f(D)$ is ε'-close to the uniform distribution over $\mathbb{Z}_N^{t-1} \times \mathbb{Z}_M$, where $\varepsilon' = \varepsilon \cdot (N^{t-1}M)^{1/2} \cdot c \log N + M/N$ and c is an absolute constant.*

Proof. By Lemma 3.6 and Corollary 3.8, $f(D)$ is $\varepsilon \cdot (N^{t-1}M)^{1/2} \cdot c \log N$-close to $f(U)$. For each $b \in \mathbb{Z}_M$, the number of $a \in \{0, 1, \dots, N-1\}$ satisfying $a \bmod M = b$ is either $\lfloor N/M \rfloor$ or $\lceil N/M \rceil$, and its difference from N/M is bounded by one. So $f(U)$ is M/N-close to the uniform distribution over $\mathbb{Z}_N^{t-1} \times \mathbb{Z}_M$, and the lemma follows. □

Small characteristic. The XOR lemma requires the distribution to be ε-biased. However, when the characteristic is small, we need to deal with the more general class of (ε, e)-biased distributions, where e is small. The following lemma states that we can extract randomness from such sources using a function f, provided that the L_1 and L_∞ norms of the Fourier transforms of certain functions $\psi \circ f$ are reasonably bounded.

Lemma 3.10. *Let A and B be finite abelian groups. Let D be an (ε, e)-biased distribution over A. Let $f : A \to B$ be a map such that for every nontrivial character $\psi \in \widehat{B}$, $\|\psi \circ f\|_1 \leq a_1$ and $\|\psi \circ f\|_\infty \leq a_\infty$. Then $f(D)$ is ε'-close to the uniform distribution over B, where $\varepsilon' = (a_1\varepsilon + a_\infty e)|B|^{1/2}$.*

Proof. Let ψ be a nontrivial character of B. By the XOR lemma, we just need to prove that

$$|\mathbb{E}[\psi(f(D))]| \leq a_1\varepsilon + a_\infty e.$$

Let E be the set of $\chi \in \widehat{A}$ such that $|\mathbb{E}[\chi(D)]| > \varepsilon$. Then $|E| \leq e$. Writing $\psi \circ f = \sum_{\chi \in \widehat{B}} \widehat{\psi \circ f}(\chi) \cdot \chi$, we have

$$\begin{aligned} |\mathbb{E}[\psi(f(D))]| &= |\mathbb{E}[(\psi \circ f)(D)]| \\ &\leq \sum_{\chi \in \widehat{A}} |\widehat{\psi \circ f}(\chi)| \cdot |\mathbb{E}[\chi(D))]| \\ &= \sum_{\chi \in \widehat{A} \setminus E} |\widehat{\psi \circ f}(\chi)| \cdot |\mathbb{E}[\chi(D))]| + \sum_{\chi \in E} |\widehat{\psi \circ f}(\chi)| \cdot |\mathbb{E}[\chi(D))]| \end{aligned}$$

$$\le \left(\sum_{\chi\in\widehat{A}\setminus E} |\widehat{\psi\circ f}(\chi)|\right)\cdot\varepsilon + |E|\cdot\max_{\chi\in E}|\widehat{\psi\circ f}(\chi)|$$
$$\le a_1\varepsilon + a_\infty e$$

as desired. □

We turn to constructing functions f with the properties required by Lemma 3.10. As a warm-up, we first give a construction based on the inner product function and error-correcting codes. Later, we give an improved construction based on *rank-metric codes.*

Construction of f based on the inner product function and error-correcting codes. We construct f by adapting the inner product function $IP(x_1,\dots,x_{2r}) = \sum_{i=1}^r x_{2i-1}x_{2i}$. Suppose $r,n>0$ are integers such that $2r\le n$. Let $\pi:\mathbb{F}_p^n\to\mathbb{F}_p^{2r}$ be a projection over $\mathbb{F}_p$. Let $C\subseteq\mathbb{F}_p^r$ be a linear code over $\mathbb{F}_p$ of dimension t and relative distance δ with a generating matrix $G=(c_{i,j})\in\mathbb{F}_p^{t\times r}$ (i.e., C is the row space of G). For $i\in[t]$, define $f_i:\mathbb{F}_p^n\to\mathbb{F}_p$ by

$$f_i(x)=\sum_{j=1}^r c_{i,j}y_{2j-1}y_{2j},\ \text{where } (y_1,\dots,y_{2r})=\pi(x)\in\mathbb{F}_p^{2r}.$$

Let $f=(f_1,\dots,f_t):\mathbb{F}_p^n\to\mathbb{F}_p^t$.

Lemma 3.11. *Let f be as above. For every nontrivial character $\psi\in\widehat{\mathbb{F}_p^t}$, we have $\left\|\widehat{\psi\circ f}\right\|_1\le p^r$ and $\left\|\widehat{\psi\circ f}\right\|_\infty\le p^{-\delta r}$.*

Proof. Note that $\psi\circ f$ may be viewed as a function in $\pi(x)\in\mathbb{F}_p^{2r}$ and hence can be written as a linear combination of the characters of $\mathbb{F}_p^{2r}$. On the other hand, the projection $\pi:\mathbb{F}_p^n\to\mathbb{F}_p^{2r}$ induces an injective group homomorphism $\iota:\widehat{\mathbb{F}_p^{2r}}\hookrightarrow\widehat{\mathbb{F}_p^n}$ via $\chi\mapsto\chi\circ\pi$. This means the support of $\widehat{\psi\circ f}$ is contained in the subgroup $\iota(\widehat{\mathbb{F}_p^{2r}})\subseteq\widehat{\mathbb{F}_p^n}$ of size p^{2r}. By the Cauchy–Schwarz inequality,

$$\left\|\widehat{\psi\circ f}\right\|_1=\sum_{\chi\in\iota(\widehat{\mathbb{F}_p^{2r}})}|\widehat{\psi\circ f}(\chi)|\le\left(\sum_{\chi\in\iota(\widehat{\mathbb{F}_p^{2r}})}|\widehat{\psi\circ f}(\chi)|^2\right)^{1/2}\cdot p^r=\left\|\widehat{\psi\circ f}\right\|_2\cdot p^r=p^r.$$

where the last equality uses the fact that $\psi\circ f$ takes values in the unit circle and hence $\left\|\widehat{\psi\circ f}\right\|_2=1$ by Parseval's identity.

We now prove the second claim, i.e., $|\widehat{\psi\circ f}(\chi)|\le p^{-\delta r}$ for $\chi\in\widehat{\mathbb{F}_p^n}$. We may also assume $\chi\in\iota(\widehat{\mathbb{F}_p^{2r}})$ since the support of $\widehat{\psi\circ f}$ is contained in $\iota(\widehat{\mathbb{F}_p^{2r}})$.

Fix a nontrivial character σ of $\mathbb{F}_p$. Then $\psi(x_1,\dots,x_t)=\sigma(\sum_{i=1}^t u_ix_i)$ for some nonzero vector $u=(u_1,\dots,u_t)\in\mathbb{F}_p^t$. Let $(v_1,\dots,v_r)=uG$, which is a nonzero codeword in C. Note that by definition,

$$(\psi\circ f)(x)=\sigma\left(\sum_{i=1}^t u_if_i(x)\right)=\sigma\left(\sum_{i=1}^r v_iy_{2i-1}y_{2i}\right),\ \text{where } (y_1,\dots,y_{2r})=\pi(x)\in\mathbb{F}_p^{2r}.$$

As $\chi \in \iota(\widehat{\mathbb{F}_p^{2r}})$, we have $\chi(x) = \sigma(\sum_{i=1}^{2r} w_i y_i)$ for some $(w_1, \dots, w_{2r}) \in \mathbb{F}_p^{2r}$, where $(y_1, \dots, y_{2r}) = \pi(x)$. Then

$$\begin{aligned}|\widehat{\psi \circ f}(\chi)| &= \left| \mathop{\mathbb{E}}_{x \in \mathbb{F}_p^n} \left[(\psi \circ f)(x) \overline{\chi(x)} \right] \right| \\ &= \left| \mathop{\mathbb{E}}_{(y_1,\dots,y_{2r}) \in \mathbb{F}_p^{2r}} \left[\sigma\left(\sum_{i=1}^{r} v_i y_{2i-1} y_{2i} \right) \cdot \overline{\sigma\left(\sum_{i=1}^{2r} w_i y_i \right)} \right] \right| \\ &= \left| \mathop{\mathbb{E}}_{(y_1,\dots,y_{2r}) \in \mathbb{F}_p^{2r}} \left[\sigma\left(\sum_{i=1}^{r} (v_i y_{2i-1} y_{2i} - w_{2i-1} y_{2i-1} - w_{2i} y_{2i}) \right) \right] \right| \\ &= \prod_{i=1}^{r} \left| \mathop{\mathbb{E}}_{y, y' \in \mathbb{F}_p} \left[\sigma\left(P_i(y, y') \right) \right] \right|\end{aligned}$$

where $P_i(y, y') := v_i y y' - w_{2i-1} y - w_{2i} y' = (v_i y - w_{2i}) y' - w_{2i-1} y$. Note that $\left| \mathbb{E}_{y, y' \in \mathbb{F}_p} \left[\sigma\left(P_i(y, y') \right) \right] \right| \leq 1/p$ whenever $v_i \neq 0$. (This holds since $\mathbb{E}_{y' \in \mathbb{F}_p} \left[\sigma\left(P_i(y, y') \right) \right]$ is zero when $v_i \neq 0$ and y is assigned a value in $\mathbb{F}_p$ different from w_{2i}/v_i, in which case $P_i(y, y')$ is a degree-1 polynomial in y'.) As C is a linear code of relative distance δ, there are at least δr indices $i \in [r]$ for which $v_i \neq 0$. It follows that $|\widehat{\psi \circ f}(\chi)| \leq p^{-\delta r}$. □

By choosing $C \subseteq \mathbb{F}_p^r$ to be an explicit asymptotically good linear code over $\mathbb{F}_p$ (e.g., an expander code), we obtain the following result.

Corollary 3.12. *There exist absolute constants $c, c' > 0$ such that the following holds. For integers $r, n > 0$ such that $2r \leq n$, there exists an explicit function $\mathbb{F}_p^n \to \mathbb{F}_p^t$, where $t \geq cr$, such that $\left\| \widehat{\psi \circ f} \right\|_1 \leq p^r$ and $L_\infty(\widehat{\psi \circ f}) \leq p^{-c' r}$ for every nontrivial character $\psi \in \widehat{\mathbb{F}_p^t}$.*

Construction of f based on rank-metric codes. The bilinear maps used in the above construction may be viewed as diagonal matrices with many nonzeros on the diagonal. We observe that the analysis only requires the matrices to have a high rank, and they do not have to be diagonal or even square matrices. This leads to the following improved construction of f based on *rank-metric codes*, which are subspaces of matrices such that every non-zero matrix in the subspace has a high rank. Here we use an optimal construction of rank-metric codes discovered independently by Delsarte [Del78] and Gabidulin [Gab85].

Definition 3.13 (Delsarte–Gabidulin codes)**.** *Let $k \leq r \leq s$ be positive integers. Fix $g_1, \dots, g_r \in \mathbb{F}_{p^s}$ that are linearly independent over $\mathbb{F}_p$. Also fix $\tau : \mathbb{F}_{p^s} \to \mathbb{F}_p^s$ that is an isomorphism of vector spaces over $\mathbb{F}_p$. For $u = (u_1, \dots, u_k) \in \mathbb{F}_{p^s}^k$, let $f_u := \sum_{i=1}^{k} u_i X^{p^{i-1}} \in \mathbb{F}_{p^s}[X]$ and $M_u := (\tau(f_u(g_1)), \dots, \tau(f_u(g_r))) \in \mathbb{F}_p^{s \times r}$, where each $\tau(f_u(g_i)) \in \mathbb{F}_p^s$ is viewed as a column vector.*

Lemma 3.14. *For any nonzero $u \in \mathbb{F}_{p^s}^k$, the matrix M_u has rank at least $r - k + 1$.*

Proof. Let $Z \subseteq \mathbb{F}_{p^s}$ be the set of roots of f_u in $\mathbb{F}_{p^s}$. We have $|Z| \leq \deg(f_u) \leq p^{k-1}$. Note that for a column vector $y = (y_1, \dots, y_r) \in \mathbb{F}_p^r$, we have $M_u y = \sum_{i=1}^{r} y_i \tau(f_u(g_i)) = \tau(f_u(\sum_{i=1}^{r} y_i g_i))$ since τ is $\mathbb{F}_p$-linear and f_u is a linearized polynomial. So $M_u y = 0$ iff $\sum_{i=1}^{r} y_i g_i \in Z$. As $g_1, \dots, g_r$ are linearly independent over $\mathbb{F}_p$, the map $y \mapsto \sum_{i=1}^{r} y_i g_i$ is injective. So there are at most $|Z| \leq p^{k-1}$ choices of y such that $M_u y = 0$. We conclude that the right kernel of M_u has size at most p^{k-1} and hence has dimension at most $k - 1$. In other words, the rank of M_u is at least $r - k + 1$. □

The map $u \mapsto M_u$ is an $\mathbb{F}_p$-linear injective map by Lemma 3.14. So $\{M_u : u \in \mathbb{F}_{p^s}^k\}$ is a linear space of dimension ks over $\mathbb{F}_p$, and any nonzero matrix in this space has rank at least $r-k+1$ by Lemma 3.14. We record this in the following corollary.

Corollary 3.15. *Let $k \le r \le s$ be positive integers. Let $t \le ks$. There exist explicit matrices $M_1, \dots, M_t \in \mathbb{F}_p^{s\times r}$ such that for any nonzero $c = (c_1, \dots, c_t) \in \mathbb{F}_p^t$, the matrix $\sum_{i=1}^t c_i M_i$ has rank at least $r-k+1$.*

Construction 3.16. *Suppose $r, n > 0$ are integers such that $2r \le n$. Let $s = n - r \ge r$. Let $k \in [r]$ and $t \in [ks]$. And let $M_1, \dots, M_t$ be as in Corollary 3.15. Identify $\mathbb{F}_p^n$ with $\mathbb{F}_p^s \times \mathbb{F}_p^r$. Then for $i \in [t]$, define $f_i : \mathbb{F}_p^n \cong \mathbb{F}_p^s \times \mathbb{F}_p^r \to \mathbb{F}_p$ to be the $\mathbb{F}_p$-bilinear map sending $(x, y) \in \mathbb{F}_p^s \times \mathbb{F}_p^r$ to $x^T M_i y$, where x and y are viewed as column vectors. Finally, let $f = (f_1, \dots, f_t) : \mathbb{F}_p^n \to \mathbb{F}_p^t$.*

Lemma 3.17. *Let f be as above. For every nontrivial character $\psi \in \widehat{\mathbb{F}_p^t}$, we have $\left\|\widehat{\psi \circ f}\right\|_1 \le p^r$ and $\left\|\widehat{\psi \circ f}\right\|_\infty \le p^{-(r-k+1)}$.*

Proof. Fix a nontrivial character σ of $\mathbb{F}_p$. Identify $\mathbb{F}_p^n$ with $\mathbb{F}_p^s \times \mathbb{F}_p^r$. Then $\psi \circ f$ has the form $(x, y) \mapsto \sigma(x^T M y)$ where $M \in \mathbb{F}^{s\times r}$ is a nontrivial linear combination of the matrices $M_1, \dots, M_t$. Let R be the rank of M. Then $r - k + 1 \le R \le \min\{s, r\} = r$ by Corollary 3.15.

Let $U, V \subseteq \mathbb{F}_p^s$ such that U is the left kernel of M and V is a complement of U. We have $\operatorname{codim}_{\mathbb{F}_p} U = \dim_{\mathbb{F}_p} V = R$. Identify $\mathbb{F}_p^s$ with $U \times V$. Identify $\widehat{\mathbb{F}_p^n}$ with $\widehat{U} \times \widehat{V} \times \widehat{\mathbb{F}_p^r}$ so that (χ_1, χ_2, χ_3) sends $(u, v, y) \in U \times V \times \mathbb{F}_p^r \cong \mathbb{F}_p^n$ to $\chi_1(u)\chi_2(v)\chi_3(y)$. Consider $\chi = (\chi_1, \chi_2, \chi_3) \in \widehat{\mathbb{F}_p^n}$. We have

$$\begin{aligned}
\widehat{\psi \circ f}(\chi) &= \mathop{\mathbb{E}}_{x \in \mathbb{F}_p^n} \left[(\psi \circ f)(x)\overline{\chi(x)}\right] \\
&= \mathop{\mathbb{E}}_{(u,v,y)\in U\times V\times \mathbb{F}_p^r} \left[\sigma((u+v)^T M y)\overline{\chi_1(u)\chi_2(v)\chi_3(y)}\right] \\
&= \mathop{\mathbb{E}}_{(u,v,y)\in U\times V\times \mathbb{F}_p^r} \left[\sigma(v^T M y)\overline{\chi_1(u)\chi_2(v)\chi_3(y)}\right] \\
&= \mathop{\mathbb{E}}_{(v,y)\in V\times \mathbb{F}_p^r} \left[\sigma(v^T M y)\overline{\chi_2(v)\chi_3(y)}\right] \mathop{\mathbb{E}}_{u\in U}\left[\overline{\chi_1(u)}\right]
\end{aligned} \tag{1}$$

Note $\mathbb{E}_{u\in U}\left[\overline{\chi_1(u)}\right] = 0$ whenever $\chi_1 \ne 1$. So the support of $\widehat{\psi \circ f}$ is contained in the subgroup $\{1\} \times \widehat{V} \times \widehat{\mathbb{F}_p^r}$ of size $p^{R+r} \le p^{2r}$. Also note that $\psi \circ f$ takes values in the unit circle of $\mathbb{C}$ and hence $\sum_{\chi\in\widehat{\mathbb{F}_p^n}} |\widehat{\psi \circ f}(\chi)|^2 = \mathbb{E}_{x\in\mathbb{F}_p^n}\left[(\psi \circ f)(x)^2\right] = 1$ by Parseval's identity. Then by the Cauchy–Schwarz inequality,

$$\left\|\widehat{\psi \circ f}\right\|_1 = \sum_{\chi\in\{1\}\times\widehat{V}\times\widehat{\mathbb{F}_p^r}} |\widehat{\psi \circ f}(\chi)| \le \left(\sum_{\chi\in\widehat{\mathbb{F}_p^n}} |\widehat{\psi \circ f}(\chi)|^2\right)^{1/2} \cdot p^r = p^r.$$

This proves the first claim.

We now prove the second claim, i.e., $|\widehat{\psi \circ f}(\chi)| \le p^{-(r-k+1)}$ for $\chi \in \widehat{\mathbb{F}_p^n}$. We may assume $\chi = (\chi_1, \chi_2, \chi_3) \in \{1\} \times \widehat{V} \times \widehat{\mathbb{F}_p^r}$ since otherwise $\widehat{\psi \circ f}(\chi) = 0$. Choose $w \in V$ such that χ_2 sends

$v \in V$ to $\sigma(v^T w)$. Then

$$\begin{aligned}
|\widehat{\psi \circ f}(\chi)| &= \left| \mathop{\mathbb{E}}_{(v,y)\in V\times\mathbb{F}_p^r} \left[\sigma(v^T My)\overline{\chi_2(v)\chi_3(y)} \right] \right| \\
&= \left| \mathop{\mathbb{E}}_{(v,y)\in V\times\mathbb{F}_p^r} \left[\sigma(v^T(My-w))\overline{\chi_3(y)} \right] \right| \\
&= \left| \mathop{\mathbb{E}}_{y\in\mathbb{F}_p^r} \left[\mathop{\mathbb{E}}_{v\in V} \left[\sigma(v^T(My-w)) \right] \cdot \overline{\chi_3(y)} \right] \right| \\
&\le \Pr_{y\in\mathbb{F}_p^r}[My = w].
\end{aligned}$$

The first equality holds by (1) and the assumption $\chi_1 = 1$. The last inequality above holds since $\mathbb{E}_{v\in V}\left[\sigma(v^T(My-w))\right] = 0$ when $My \neq w$. Finally, note $\Pr_{y\in\mathbb{F}_p^r}[My = w]$ is either zero or p^{-R}, depending on whether $My = w$ has a solution. In either case, we have $|\widehat{\psi \circ f}(\chi)| \le p^{-R} \le p^{-(r-k+1)}$. □

Assuming that n is sufficiently large and ε is not too much larger than $p^{-n/2}$, the construction above gives explicit deterministic extractors that extract a constant fraction of min-entropy from (ε, e)-biased sources over $\mathbb{F}_p^n$, as stated by the following theorem.

Theorem 3.18. *Let $n, t, e \in \mathbb{N}^+$ and $d, \varepsilon' > 0$ such that $n \ge c$ and $t \log p \le c^{-1} n \log p - 2\log(de/\varepsilon')$, where $c > 0$ is a large enough absolute constant. Let $\varepsilon = dp^{-n/2}$. Then there exists an explicit deterministic ε'-extractor $f : \mathbb{F}_p^n \to \mathbb{F}_p^t$ for (ε, e)-biased sources distributed over $\mathbb{F}_p^n$.*

Proof. Choose $r = n/4$, $s = n - r$, and $k = r/2$. (For ease of readability, we omit floor and ceiling functions.) Note $t \le ks$ assuming c is large enough. Let $f : \mathbb{F}_p^n \to \mathbb{F}_p^t$ be as in Construction 3.16. Then $\left\|\widehat{\psi \circ f}\right\|_1 \le p^r$ and $\left\|\widehat{\psi \circ f}\right\|_\infty \le p^{-r/2}$ by Lemma 3.17. By Lemma 3.10, for any (ε, e)-biased distribution D over $\mathbb{F}_p^n$, $f(D)$ is ε''-close to the uniform distribution over $\mathbb{F}_p^t$, where $\varepsilon'' = (p^r dp^{-n/2} + p^{-r/2}e) \cdot p^{t/2}$. As c is large enough, the conditions in the theorem imply $p^r dp^{-n/2} \cdot p^{t/2} \le \varepsilon'/2$ and $p^{-r/2}e \cdot p^{t/2} \le \varepsilon'/2$. So $\varepsilon'' \le \varepsilon'$ and hence $f(D)$ is ε'-close to the uniform distribution. □

We also obtain an explicit construction of deterministic extractors that extract most min-entropy from very dense affine sources, and more generally, $(0, e)$-biased sources. It is interesting to compare and constant the theorem below with our construction in Section 10, which requires a large field size but works for arbitrary min-entropy, and is also based on very different ideas.

Theorem 3.19. *Let $n, t, e \in \mathbb{N}^+$ and $\varepsilon \in (0, 1)$ such that $t \le n - 3 - 2\log_p(e/\varepsilon)$. Then there exists an explicit deterministic ε-extractor $f : \mathbb{F}_p^n \to \mathbb{F}_p^t$ for $(0, e)$-biased sources, and in particular, for affine sources of codimension at most $\log_p e$.*

Proof. Note that $n \ge t + 3 \ge 4$. Choose $r = \lfloor n/2 \rfloor \ge (n-1)/2$, $s = n - r$, and $k = 2 \le r$. Note $t \le ks$ since $t \le n$ and $s \ge n/2$. Let $f : \mathbb{F}_p^n \to \mathbb{F}_p^t$ be as in Construction 3.16. Then $\left\|\widehat{\psi \circ f}\right\|_\infty \le p^{-(r-k+1)}$ by Lemma 3.17. By Lemma 3.10, for any $(0, e)$-biased distribution D over $\mathbb{F}_p^n$, $f(D)$ is ε'-close to the uniform distribution over $\mathbb{F}_p^t$, where $\varepsilon' = p^{-(r-k+1)} \cdot e \cdot p^{t/2}$. Finally, note that $\varepsilon' \le \varepsilon$ by the choice of t. □

Extracting most min-entropy from strongly (ε, e)-biased sources. We now construct deterministic extractors that extract most min-entropy from strongly (ε, e)-biased sources.

The following lemma allows us to reduce to the case of affine sources.

Lemma 3.20. *Let $\pi : \mathbb{F}_p^n \to \mathbb{F}_p^m$ be a surjective $\mathbb{F}_p$-linear map. Let D be a strongly (ε, e)-biased distribution over $\mathbb{F}_p^n$. Then $\pi(D)$ is ε'-close to a convex combination of affine sources of codimension at most $\log_p e$ in $\mathbb{F}_p^m$, where $\varepsilon' = 2(p^{m/2} \cdot e \cdot \varepsilon)^{1/2}$.*

Proof. Let A be a subspace of $\mathbb{F}_p^n$ of size at most e and B be a complement of A such that by identifying $\mathbb{F}_p^n$ with $A \times B$ and $\widehat{\mathbb{F}_p^n}$ with $\widehat{A} \times \widehat{B}$, every nontrivial character χ of $\widehat{\mathbb{F}_p^n}$ satisfying $|\mathbb{E}[\chi(D)]| > \varepsilon$ is in the subgroup $\widehat{A} \times \{1\}$. This is possible as D is strongly (ε, e)-biased. Identify A and B with the subgroups $A \times \{0\}$ and $\{0\} \times B$ of $A \times B$ respectively. Then the codimension of B in $\mathbb{F}_p^n$ is at most $\log_p e$. Let D_1 and D_2 be the marginal distributions of D over A and B respectively.

Let $\varepsilon_0 > 0$, whose value is determined later. By Corollary 3.4, for x sampled from D_1, with probability at least $1-\varepsilon_0$, the conditional distribution $D_2|_{D_1=x}$ is $|A|\varepsilon/\varepsilon_0$-biased. Fix $x \in \operatorname{supp}(D_1)$ such that $D_2|_{D_1=x}$ is $|A|\varepsilon/\varepsilon_0$-biased. We claim that $\pi(D|_{D_1=x})$ is ε_1-close to an affine source of codimension at most $\log_p e$, where $\varepsilon_1 = p^{m/2}|A|\varepsilon/\varepsilon_0$. Note that $\pi(D|_{D_1=x}) = \pi(x) + \pi(D_2|_{D_1=x})$. So to prove the claim, we just need to show that $\pi(D_2|_{D_1=x})$ is ε_1-close to an affine source of codimension at most $\log_p e$.

Let $H = \pi(B)$. Then $\pi|_B : B \to H$ is a surjective linear map. So the map $\chi \mapsto \chi \circ \pi$ sends a nontrivial character of H to a nontrivial character of B. As $D_2|_{D_1=x}$ is an $|A|\varepsilon/\varepsilon_0$-biased distribution over B, we see that $\pi(D_2|_{D_1=x})$ is an $|A|\varepsilon/\varepsilon_0$-biased distribution over H. By the XOR lemma (Lemma 3.5), $\pi(D_2|_{D_1=x})$ is $|H|^{1/2}|A|\varepsilon/\varepsilon_0$-close to the uniform distribution over H. Note

$$\begin{aligned}\dim H = \dim B - \dim(B \cap \ker \pi) &\geq \dim B - \dim(\ker \pi) \\ &= \dim B - (n-m) = m - \operatorname{codim} B.\end{aligned}$$

So the codimension of H in $\mathbb{F}_p^m$ is at most $\operatorname{codim} B \leq \log_p e$. Therefore, the distribution $\pi(D_2|_{D_1=x})$ is $|H|^{1/2}|A|\varepsilon/\varepsilon_0$-close to an affine source of codimension at most $\log_p e$. This proves the above claim as $|H|^{1/2}|A|\varepsilon/\varepsilon_0 \leq p^{m/2}|A|\varepsilon/\varepsilon_0 = \varepsilon_1$.

We have shown that for x sampled from D_1, with probability at least $1-\varepsilon_0$, the distribution $D|_{D_1=x}$ is ε_1-close to an affine source of codimension at most $\log_p e$. It follows that D is $(\varepsilon_0+\varepsilon_1)$-close to a convex combination of affine sources of codimension at most $\log_p e$, where $\varepsilon_1 = p^{m/2}|A|\varepsilon/\varepsilon_0$. Finally, choose $\varepsilon_0 = (p^{m/2}|A|\varepsilon)^{1/2}$ so that

$$\varepsilon_0 + \varepsilon_1 = 2\varepsilon_0 = 2(p^{m/2} \cdot |A| \cdot \varepsilon)^{1/2} \leq 2(p^{m/2} \cdot e \cdot \varepsilon)^{1/2}. \qquad \square$$

Theorem 3.21. *Let n, t, e be positive integers and $\varepsilon, \varepsilon' \in (0,1)$. Let $n' = \min\{\lfloor 2\log_p(1/\varepsilon) - 2\log_p(16e/\varepsilon'^2)\rfloor, n\}$. Suppose $t \leq n' - 3 - 2\log_p(2e/\varepsilon')$. Then there exists an explicit ε'-extractor $\mathsf{Ext} : \mathbb{F}_p^n \to \mathbb{F}_p^t$ for strongly (ε, e)-biased sources.*

Proof. We construct the extractor Ext as follows.

- Let $\pi : \mathbb{F}_p^n \to \mathbb{F}_p^{n'}$ be an explicit linear surjective map.
- Let $f : \mathbb{F}_p^{n'} \to \mathbb{F}_p^t$ be an explicit deterministic $\varepsilon'/2$-extractor for affine sources of codimension at most $\log_p e$.

- Finally, let $\mathsf{Ext} = f \circ \pi$.

Here the existence of f is guaranteed by Theorem 3.19 as $t \leq n' - 3 - 2\log_p(2e/\varepsilon')$. We claim that $\mathsf{Ext} : \mathbb{F}_p^n \to \mathbb{F}_p^t$ is a deterministic ε'-extractor for strongly (ε, e)-biased sources.

Consider an arbitrary strongly (ε, e)-biased source D. By the choice of n', we have $2(p^{n'/2} \cdot e \cdot \varepsilon)^{1/2} \leq \varepsilon'/2$. So by Lemma 3.20, $\pi(D)$ is $\varepsilon'/2$-close to a convex combination of affine sources of codimension at most $\log_p e$. Then by the affine extractor property of f, we know $\pi(D)$ is ε'-close to the uniform distribution. This proves the claim. □

Remark 3.22. By Proposition 3.2, an (ε, e)-biased source over $\mathbb{F}_p^n$ is ε'-close to a source of min-entropy at least $\min\{2\log(1/\varepsilon), n\log p - \log e\} - \log(2/\varepsilon')$. The output bit-length $t\log p$ in Theorem 3.21 matches this bound up to an additive term $O(\log e + \log(1/\varepsilon') + \log p)$.

Remark 3.23. As (ε, e)-biased sources over $\mathbb{F}_p^n$ are strongly (ε, p^e)-biased, Theorem 3.21 also gives deterministic extractors for (ε, e)-biased sources at the cost of replacing e with p^e in the theorem statement. It is interesting to ask for deterministic extractors, explicit or not, that extract most min-entropy from (ε, e)-biased sources (up to the lower bound in Proposition 3.2) without this exponential blow-up. We note that the extractors in Theorem 3.18 could extract a constant fraction of the min-entropy from (ε, e)-biased sources assuming that ε is small enough.

4 Preliminaries on Algebraic Geometry

In this section, we introduce preliminaries and notations on algebraic geometry. One can also refer to a standard text, e.g., [Sha94, Vak22].

4.1 Terminology

All rings in this paper are commutative rings with unity. A proper ideal I of a ring R is *prime* if $ab \in I$ implies $a \in I$ or $b \in I$ for $a, b \in R$. This is equivalent to the condition that R/I is an integral domain. An ideal I of R is *radical* if $a^m \in I$ implies $a \in I$ for every $a \in R$ and $m \in \mathbb{N}^+$.

Affine varieties. Let $\mathbb{F}$ be a field and let $\overline{\mathbb{F}}$ be its algebraic closure. For $n \in \mathbb{N}$, the *affine n-space* $\mathbb{A}_{\mathbb{F}}^n$ *over* $\mathbb{F}$ is the set $\overline{\mathbb{F}}^n$ equipped with the *Zariski topology*, defined as follows. A subset $U \subseteq \mathbb{A}_{\mathbb{F}}^n$ is *closed* if it is the set of common zeros of a set of polynomials in $\mathbb{F}[X_1, \ldots, X_n]$. And $U \subseteq \mathbb{A}_{\mathbb{F}}^n$ is *open* if its complement is closed. A closed subset $V \subseteq \mathbb{A}_{\mathbb{F}}^n$ is said to be an *affine variety over* $\mathbb{F}$, and its elements are called *points* of V.[2] We say V is *defined by* a set $S \subseteq \mathbb{F}[X_1, \ldots, X_n]$ if it is the set of common zeros of the polynomials in S.

We often write $\mathbb{A}^n$ instead of $\mathbb{A}_{\mathbb{F}}^n$ when $\mathbb{F}$ is algebraically closed and clear from the context.

A point $a \in \mathbb{A}_{\mathbb{F}}^n$ is said to be a *rational point* if the coordinates of a are all in $\mathbb{F}$. For an affine variety V over $\mathbb{F}$, denote by $V(\mathbb{F})$ the set of rational points in V. Each rational point is closed (as a set containing the point itself) in the Zariski topology.

[2] When $\mathbb{F}$ is not algebraically closed, affine varieties (and affine spaces) are often defined in a different way such that each point of an affine variety is not a single point in $\overline{\mathbb{F}}^n$, but an orbit consisting of all the conjugates of a point in $\overline{\mathbb{F}}^n$ under the natural action of the automorphism group of $\overline{\mathbb{F}}$ over $\mathbb{F}$. Our definition instead follows that in [Dvi12], which suffices for us and is more elementary. One can switch between these two definitions by splitting the orbits into points in $\overline{\mathbb{F}}^n$ and vice versa. One notable difference between the two definitions is that, in our definition, points are not necessarily closed in the Zariski topology over $\mathbb{F}$ when the field $\mathbb{F}$ is not algebraically closed, while they are always closed in the other definition.

For an affine variety $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ over $\mathbb{F}$ defined by $S \subseteq \mathbb{F}[X_1, \dots, X_n]$, and an extension field $\mathbb{K}$ of $\mathbb{F}$, denote by $V_{\mathbb{K}}$ the affine variety $V' \subseteq \mathbb{A}^n_{\mathbb{K}}$ over $\mathbb{K}$ defined by S.

Remark. Note that under our definitions, $V_{\overline{\mathbb{F}}}$ is the same as V as subsets of $\overline{\mathbb{F}}^n$.

For two affine varieties $V, V' \subseteq \mathbb{A}^n_{\mathbb{F}}$ over a field $\mathbb{F}$, we say V is an *affine subvariety* of V' if $V \subseteq V'$. We say V is the affine subvariety of V' *defined by* $S \subseteq \mathbb{F}[X_1, \dots, X_n]$ if it is the set of common zeros in V' of the polynomials in S.

Irreducibility and absolute irreducibility. An affine variety V over a field $\mathbb{F}$ is *irreducible* if it is nonempty and cannot be written as the union of finitely many proper affine subvarieties over $\mathbb{F}$. Otherwise, we say V is *reducible.* A subvariety V_0 of an affine variety V is an *irreducible component* of V if V_0 is irreducible and maximal (with respect to set inclusion) for this property. Every affine variety can be uniquely represented as the union of finitely many irreducible components.

An affine variety V over $\mathbb{F}$ is *absolutely irreducible* if $V_{\overline{\mathbb{F}}}$ is irreducible. By definition, the Zariski topology over $\overline{\mathbb{F}}$ is finer than that over $\mathbb{F}$, i.e., if a set is closed over $\mathbb{F}$, then it is also closed over $\overline{\mathbb{F}}$. So absolute irreducibility implies irreducibility.

Let $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ be an irreducible variety over $\mathbb{F}$. The automorphism group $\mathrm{Aut}(\overline{\mathbb{F}}/\mathbb{F})$ of $\overline{\mathbb{F}}$ over $\mathbb{F}$ acts naturally on $V_{\overline{\mathbb{F}}}$ such that $\tau \in \mathrm{Aut}(\overline{\mathbb{F}}/\mathbb{F})$ sends a point $a = (a_1, \dots, a_n) \in V_{\overline{\mathbb{F}}}$ to $(\tau(a_1), \dots, \tau(a_n))$. Every $\tau \in \mathrm{Aut}(\overline{\mathbb{F}}/\mathbb{F})$ permutes the irreducible components of $V_{\overline{\mathbb{F}}}$. When $\mathbb{F} = \mathbb{F}_q$, the Frobenius automorphism $\sigma : x \mapsto x^q$ of $\overline{\mathbb{F}}_q$ over $\mathbb{F}_q$ generates a cyclic group that acts transitively on the set of irreducible components of $V_{\overline{\mathbb{F}}}$. In particular, if V is not absolutely irreducible, then σ sends each irreducible component of $V_{\overline{\mathbb{F}}}$ to a different irreducible component.

The ideal-variety correspondence. For an ideal I of $\mathbb{F}[X_1, \dots, X_n]$, denote by $V(I)$ the affine subvariety of $\mathbb{A}^n_{\mathbb{F}}$ over $\mathbb{F}$ defined by I. Define $V(f_1, \dots, f_k) = V(\langle f_1, \dots, f_k \rangle)$ for $f_1, \dots, f_k \in \mathbb{F}[X_1, \dots, X_n]$. For an affine variety $V \subseteq \mathbb{A}^n$ over $\mathbb{F}$, denote by $I(V)$ the ideal of $\mathbb{F}[X_1, \dots, X_n]$ consisting of all the polynomials vanishing on V. The ideal $I(V)$, and in fact every ideal of $\mathbb{F}[X_1, \dots, X_n]$, has a finite generating set by *Hilbert's Basis Theorem.*

The map $V \mapsto I(V)$ is an inclusion-reversing one-to-one correspondence between the affine subvarieties of $\mathbb{A}^n_{\mathbb{F}}$ over $\mathbb{F}$ and the radical ideals of $\mathbb{F}[X_1, \dots, X_n]$, with the inverse map $I \mapsto V(I)$. An affine variety V is irreducible iff $I(V)$ is a prime ideal.

For an affine variety $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ over $\mathbb{F}$, define

$$\mathbb{F}[V] := \mathbb{F}[X_1, \dots, X_n]/I(V),$$

which is called the *coordinate ring* of V. When V is irreducible, the ideal $I(V)$ is prime and $\mathbb{F}[V]$ is an integral domain. In this case, define the *function field* of V, denoted by $\mathbb{F}(V)$, to be the field of fractions of $\mathbb{F}[V]$.

Dimension. The *dimension* of an irreducible affine variety V over a field $\mathbb{F}$ is defined to be the largest integer m such that there exists a chain of irreducible affine subvarieties $\emptyset \subsetneq V_0 \subsetneq V_1 \subsetneq \cdots \subsetneq V_m = V$ over $\mathbb{F}$. More generally, the dimension of a nonempty affine variety V, denoted by $\dim V$, is the maximal dimension of its irreducible components. We have $\dim V = \dim V_{\overline{\mathbb{F}}}$ for a nonempty affine variety V over $\mathbb{F}$.

The dimension of an irreducible affine variety V over $\mathbb{F}$ also equals the *transcendence degree* of $\mathbb{F}(V)/\mathbb{F}$, i.e., the largest cardinality of an algebraically independent subset of $\mathbb{F}(V)$ over $\mathbb{F}$. In particular, $\dim \mathbb{A}^n_{\mathbb{F}} = n$ as its function field $\mathbb{F}(X_1, \dots, X_n)$ has transcendence degree n over $\mathbb{F}$.

A nonempty affine variety is *equidimensional* if its irreducible components have the same dimension. Equidimensional affine varieties of dimension one are also called *(affine) curves.*

The following fact is a geometric version of *Krull's principal ideal theorem.* See, e.g., [Vak22, §12.3].

Lemma 4.1. *Let $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ be an irreducible affine variety over a field $\mathbb{F}$. Let $f \in \mathbb{F}[X_1, \dots, X_n]$ such that f does not vanish identically on V. Then either $V \cap V(f) = \emptyset$ or $V \cap V(f)$ is equidimensional of dimension* $\dim V - 1$.

Degree. Let $\mathbb{F}$ be an algebraically closed field. For an irreducible affine variety $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ over $\mathbb{F}$, the *degree* of V in $\mathbb{A}^n_{\mathbb{F}}$, denoted by $\deg V$, is the number of intersections of V with an affine subspace of $\mathbb{A}^n_{\mathbb{F}}$ of codimension $\dim V$ in general position. Following [HS80, Hei83], for an affine variety $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ over $\mathbb{F}$ with irreducible components $V_1, \dots, V_k$, we let $\deg V := \sum_{i=1}^k \deg V_i$. For an affine variety V over an arbitrary field $\mathbb{F}$, let $\deg V := \deg V_{\overline{\mathbb{F}}}$.

For a nonzero polynomial $f \in \mathbb{F}[X_1, \dots, X_n]$, we have $\deg(V(f)) \leq \deg f$ (and the equality holds if f is squarefree).

The following version of *Bézout's inequality* is very useful for us.

Lemma 4.2 (Bézout's inequality [HS80, Hei83])**.** *Let $V, V' \subseteq \mathbb{A}^n_{\mathbb{F}}$ be affine varieties over a field $\mathbb{F}$. Then*

$$\deg(V \cap V') \leq \deg V \cdot \deg V'.$$

Morphisms between affine varieties over a field $\mathbb{F}$. Let $\mathbb{F}$ be a field. Let $f_1, \dots, f_m \in \mathbb{F}[X_1, \dots, X_n]$, which define a map $f : \mathbb{A}^n_{\mathbb{F}} \to \mathbb{A}^m_{\mathbb{F}}$ sending $a \in \mathbb{A}^n_{\mathbb{F}}$ to $f(a) = (f_1(a), \dots, f_m(a)) \in \mathbb{A}^m_{\mathbb{F}}$. Now suppose $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ and $V' \subseteq \mathbb{A}^m_{\mathbb{F}}$ are affine varieties over $\mathbb{F}$ such that $f(V) \subseteq V'$. Then the map f restricts to a map $\varphi : V \to V'$. It is associated with an $\mathbb{F}$-algebra homomorphism

$$\varphi^\sharp : \mathbb{F}[V'] = \mathbb{F}[Y_1, \dots, Y_m]/I(V') \to \mathbb{F}[V] = \mathbb{F}[X_1, \dots, X_n]/I(V),$$

which sends $Y_i + I(V')$ to $f_i(X_1, \dots, X_n) + I(V)$ for $i \in [m]$. Note that $\varphi^\sharp$ simply sends a function on V' to its composition with φ, which is a function on V.

We say the pair $(\varphi, \varphi^\sharp)$ is a *morphism* from V to V' (over $\mathbb{F}$) and it is *defined by* $f_1, \dots, f_m$. For simplicity, we usually suppress $\varphi^\sharp$ and denote the morphism by φ when there is no confusion.

A morphism between affine spaces is also called a *polynomial map* in this paper. Let $\varphi : \mathbb{A}^n_{\mathbb{F}} \to \mathbb{A}^m_{\mathbb{F}}$ be a polynomial map. If $V \subseteq \mathbb{A}^m_{\mathbb{F}}$ is an affine variety defined by a set S of polynomials, then $\varphi^{-1}(V) \subseteq \mathbb{A}^n_{\mathbb{F}}$ is the affine variety defined by $\varphi^\sharp(S)$. So the preimage of a closed (resp. open) set under a polynomial map is closed (resp. open).

Zariski closure and dominant morphisms. For a set $S \subseteq \mathbb{A}^n_{\mathbb{F}}$, the (*Zariski*-)closure of S, denoted by $\overline{S}$, is the smallest closed set containing S, i.e., the intersection of all affine subvarieties of $\mathbb{A}^n_{\mathbb{F}}$ that contain S. We say S is *dense* in V if $\overline{S} = V$.

Suppose $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ is an affine variety over $\mathbb{F}$ and $\varphi : V \to \mathbb{A}^m_{\mathbb{F}}$ is a morphism defined by polynomials $f_1, \dots, f_m \in \mathbb{F}[X_1, \dots, X_n]$. The image $\varphi(V)$ is not necessarily a closed set in $\mathbb{A}^n_{\mathbb{F}}$. To understand the closure $\overline{\varphi(V)}$, note that a polynomial $P \in \mathbb{F}[Y_1, \dots, Y_m]$ vanishes on $\varphi(V)$ (or equivalently, on $\overline{\varphi(V)}$) iff the composition of P with φ vanishes on V, i.e., $\varphi^\sharp(P) = 0$. So the ideal of $\mathbb{F}[Y_1, \dots, Y_n]$ defining $\overline{\varphi(V)}$ is precisely the kernel of $\varphi^\sharp : \mathbb{F}[Y_1, \dots, Y_n] \to \mathbb{F}[V]$. Let

$r_i = X_i + I(V) \in \mathbb{F}[X_1, \dots, X_n]/I(V) = \mathbb{F}[V]$ for $i \in [n]$. Then the coordinate ring $\mathbb{F}[\overline{\varphi(V)}] = \mathbb{F}[Y_1, \dots, Y_m]/I\left(\overline{\varphi(V)}\right)$ of $\overline{\varphi(V)}$ may be identified with $\mathbb{F}[f_1(r_1, \dots, r_n), \dots, f_m(r_1, \dots, r_n)] \subseteq \mathbb{F}[V]$ via $Y_i + I\left(\overline{\varphi(V)}\right) \mapsto f_i(r_1, \dots, r_n)$.

We say a morphism $\varphi : V \to V'$ between affine varieties is *dominant* if $\overline{\varphi(V)} = V'$. If $\varphi : V \to V'$ is a dominant morphism between affine varieties and V is irreducible, then V' is also irreducible.

We also need the following lemma.

Lemma 4.3. *Let $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ be an irreducible affine variety over a field $\mathbb{F}$. Let $f_1, \dots, f_s, g_1, \dots, g_t \in \mathbb{F}[X_1, \dots, X_n]$. The polynomials $f_1, \dots, f_s$ (resp. $f_1, \dots, f_s, g_1, \dots, g_t$) define a polynomial map $\pi_1 : \mathbb{A}^n_{\mathbb{F}} \to \mathbb{A}^s_{\mathbb{F}}$ (resp. $\pi_2 : \mathbb{A}^n_{\mathbb{F}} \to \mathbb{A}^{s+t}_{\mathbb{F}}$). Then there exist distinct $i_1, \dots, i_k \in [t]$, where $k = \dim \overline{\pi_2(V)} - \dim \overline{\pi_1(V)}$, such that the polynomial map $\pi : \mathbb{A}^n_{\mathbb{F}} \to \mathbb{A}^{s+k}_{\mathbb{F}}$ defined by $f_1, \dots, f_s, g_{i_1}, \dots, g_{i_k}$ satisfies $\dim \overline{\pi(V)} = \dim \overline{\pi_2(V)} = \dim \overline{\pi_1(V)} + k$.*

Proof. Let $r_i = X_i + I(V) \in \mathbb{F}[X_1, \dots, X_n]/I(V) = \mathbb{F}[V]$ for $i \in [n]$. Let $\bar{f}_i = f(r_1, \dots, r_n)$ for $i \in [s]$ and $\bar{g}_i = g(r_1, \dots, r_n)$ for $i \in [t]$. The function fields of $\overline{\pi_1(V)}$ and $\overline{\pi_2(V)}$ are $\mathbb{K}_1 := \mathbb{F}(\bar{f}_1, \dots, \bar{f}_s)$ and $\mathbb{K}_2 := \mathbb{F}(\bar{f}_1, \dots, \bar{f}_s, \bar{g}_1, \dots, \bar{g}_t)$ respectively. So $\bar{g}_1, \dots, \bar{g}_t$ is a generating set of $\mathbb{K}_2/\mathbb{K}_1$. The lemma then follows from the fact in field theory that a generating set of a field extension $\mathbb{K}_2/\mathbb{K}_1$ always contains a subset that is a transcendence basis of $\mathbb{K}_2/\mathbb{K}_1$ (see [Lan02, Chapter VIII, Theorem 1.1]). □

Fibers of morphisms. Let $\varphi : V \to V'$ be a morphism between affine varieties over $\mathbb{F}$. For $b \in V'(\mathbb{F})$, the preimage $\varphi^{-1}(b) = \{a \in V : \varphi(a) = b\}$ is an affine subvariety of V over $\mathbb{F}$. We call $\varphi^{-1}(b)$ the *fiber of φ over b*, or the *fiber of V over b* if φ is clear from the context.

The following theorem, known as the *fiber dimension theorem*, relates the dimension of a general fiber of a dominant morphism $\varphi : V \to V'$ between irreducible affine varieties V, V' to the dimension of V and that of V'.

Theorem 4.4 (Fiber dimension theorem)**.** *Suppose $\varphi : V \to V'$ is a dominant morphism between irreducible affine varieties over an algebraically closed field $\mathbb{F}$. Then for every $b \in \varphi(V)$ and every irreducible component Z of $\varphi^{-1}(b)$, it holds that*

$$\dim Z \geq \dim V - \dim V'.$$

Moreover, there exists $U \subseteq \varphi(V)$ such that U is a dense open subset of V' and $\dim \varphi^{-1}(b) = \dim V - \dim V'$ holds for all $b \in U$.

See, e.g., [Sha94, §I.6.3, Theorem 7] for a proof. We remark that the above version of the fiber dimension theorem can be generalized in several ways, but this version suffices for us.

We also need the notion of generic fibers. Let $\varphi : V \to V'$ be a dominant morphism defined by $f_1, \dots, f_m \in \mathbb{F}[X_1, \dots, X_n]$ between irreducible affine varieties $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ and $V' \subseteq \mathbb{A}^m_{\mathbb{F}}$ over $\mathbb{F}$. Then $\mathbb{F}[V'] = \mathbb{F}[Y_1, \dots, Y_m]/I(V')$ is identified with a subring of $\mathbb{F}[V]$ under $\varphi^\sharp$, so that $Y_i + I(V')$ is identified with $\bar{f}_i := f_i + I(V)$. The *generic fiber* V_φ of φ is then the affine subvariety of $V_{\mathbb{F}(V')}$ defined by $f_1 - \bar{f}_1, \dots, f_m - \bar{f}_m \in \mathbb{F}(V')[X_1, \dots, X_n]$. Its known from transcendence theory that V_φ is irreducible of dimension $\dim V - \dim V'$. Thus, the second claim in Theorem 4.4 states that the dimension of the irreducible components of a general fiber equals that of the generic fiber.

Finite morphisms. We say a morphism $\varphi : V \to V'$ between affine varieties V and V' over $\mathbb{F}$ is a *finite* morphism if $\mathbb{F}[V]$ is finitely generated as a module over its subring $\varphi^\sharp(\mathbb{F}[V'])$. The image of a closed set under a finite morphism is closed. In particular, a finite morphism is surjective if it is dominant. Fibers of finite morphisms are finite sets.

4.2 Further Results

We list some further results in algebraic geometry, which are used in later sections.

Estimates for the number of rational points over $\mathbb{F}_q$. We start with the following elementary upper bound on the number of rational points of an affine variety over a finite field $\mathbb{F}_q$, which can be derived from Bézout's inequality [HS80, CM06].

Lemma 4.5 ([HS80, Proposition 2.3])**.** *Let $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ be an affine variety of dimension k and degree d over a field $\mathbb{F}$. Let $S \subseteq \mathbb{F}$ be a finite set. Then $|V \cap S^n| \le d|S|^k$. In particular, if $\mathbb{F} = \mathbb{F}_q$, then $|V(\mathbb{F}_q)| \le dq^k$.*

If V is absolutely irreducible and q is sufficiently large, one can do better than Lemma 4.5 and show that $|V(\mathbb{F}_q)|$ is close to $q^{\dim V}$ using the *Lang–Weil bound.* We need the following effective version of this bound.

Theorem 4.6 (Effective Lang–Weil bound)**.** *Let $V \subseteq \mathbb{A}^n_{\mathbb{F}_q}$ be an absolutely irreducible affine variety over $\mathbb{F}_q$ of dimension k and degree d. Then*

$$|V(\mathbb{F}_q) - q^k| < (d-1)(d-2)q^{k-1/2} + 5d^{13/3}q^{k-1}.$$

In particular, we have $|V(\mathbb{F}_q)| \ge q^k/2$ if $q \ge 20d^5$.

Theorem 4.6 was proved by Cafure and Matera as [CM06, Theorem 7.1] with an extra condition that $q > 2(k+1)d^2$. However, a more careful analysis shows that this condition can be removed. This was confirmed to us in [Mat22]. See Appendix A for more details. The fact that q does not need to depend on k is crucial to making our required field size independent of $\dim V$, where V is an affine variety that defines an (n, k, d) algebraic source over $\mathbb{F}_q$.

Bombieri's estimate for exponential sums. We also need Bombieri's estimate for exponential sums over rational points of curves over $\mathbb{F}_q$.

Theorem 4.7 ([Bom66, Theorem 6])**.** *Let $C \subseteq \mathbb{A}^n_{\mathbb{F}_q}$ be an affine curve of degree d_1 over a finite field $\mathbb{F}_q$ of characteristic p. Let $\sigma : \mathbb{F}_p \to \mathbb{C}^\times$ be the character $x \mapsto e^{2\pi i x/p}$ of $\mathbb{F}_p$. Suppose $f \in \mathbb{F}_q[X_1, \ldots, X_n]$ is a polynomial of degree d_2 such that for any $g \in \overline{\mathbb{F}}_q[X_1, \ldots, X_n]$ and any irreducible component C_0 of C, the function $f - (g^p - g)$ does not vanish identically on C_0. Then*

$$\left| \sum_{x \in C(\mathbb{F}_q)} (\sigma \circ \mathrm{Tr} \circ f)(x) \right| \le (d_1^2 + 2d_1 d_2 - 3d_1)q^{1/2} + d_1^2.$$

where Tr *denotes the trace map from $\mathbb{F}_q$ to $\mathbb{F}_p$.*

Noether normalization. The classical Noether normalization lemma proved by Noether [Noe26] states that an affine variety V of dimension k over an infinite field $\mathbb{F}$ admits a finite morphism $\varphi : V \to \mathbb{A}^k_{\mathbb{F}}$. Moreover, φ may be chosen to be a linear map. We give the following quantitative version of this result, which states that the coefficients that specify the linear map can be chosen from a finite subset $S \subseteq \mathbb{F}$ provided that S is large enough.

Lemma 4.8 (Noether normalization)**.** *Let $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ be an affine variety of dimension k and degree d over a field $\mathbb{F}$. Suppose S is a finite subset of $\mathbb{F}$ of size greater than d. Then there exists a polynomial map $\varphi : \mathbb{A}^n_{\mathbb{F}} \to \mathbb{A}^k_{\mathbb{F}}$ defined by linear polynomials $\ell_i = \sum_{j=1}^n c_{i,j} X_i \in \mathbb{F}[X_1, \dots, X_n]$ with coefficients $c_{i,1}, \dots, c_{i,n} \in S$ for $i = 1, \dots, k$ such that $\varphi|_V : V \to \mathbb{A}^k_{\mathbb{F}}$ is a finite morphism.*

For convenience, we also prove the following lemma, which guarantees the existence of linear polynomials achieving simultaneous Noether normalization for two affine varieties.

Lemma 4.9. *Let $\mathbb{K}_1$ and $\mathbb{K}_2$ be extension fields of a field $\mathbb{F}$. For $i = 1, 2$, let $V_i \subseteq \mathbb{A}^n_{\mathbb{K}_i}$ be an affine variety of dimension k_i and degree d_i over $\mathbb{K}_i$. Suppose S is a finite subset of $\mathbb{F}$ of size greater than $d_1 + d_2$. Then there exist linear polynomials $\ell_1, \dots, \ell_{\max\{k_1,k_2\}} \in \mathbb{F}[X_1, \dots, X_n]$ with coefficients in S such that the morphism $V_i \to \mathbb{A}^{k_i}_{\mathbb{K}_i}$ defined by $\ell_1, \dots, \ell_{k_i}$ is finite for $i = 1, 2$.*

See Appendix A for the proofs of Lemma 4.8 and Lemma 4.9.

Effective fiber dimension theorem. We also need an effective version of the fiber dimension theorem. To suit our needs, we first formulate the theorem in the following general form. Recall that for $h_1, \dots, h_s \in \mathbb{F}[X_1, \dots, X_n]$, we denote by $\mathcal{L}_{h_1,\dots,h_s,\mathbb{F}}$ the linear span of $h_1, \dots, h_s$ and 1 over $\mathbb{F}$.

Theorem 4.10 (Effective fiber dimension theorem – general form)**.** *Let $V \subseteq \mathbb{A}^n$ be an irreducible affine variety of dimension k over an algebraically closed field $\mathbb{F}$. Let $h_1, \dots, h_s \in \mathbb{F}[X_1, \dots, X_n]$ with $\deg h_1 \geq \cdots \geq \deg h_s$. Let $f_1, \dots, f_m \in \mathcal{L}_{h_1,\dots,h_s,\mathbb{F}}$, which define a polynomial map $f : \mathbb{A}^n \to \mathbb{A}^m$. Let $k' = \dim \overline{f(V)}$.*

Let $j_1, \dots, j_{k'} \in [m]$ such that the morphism $f' : V \to \mathbb{A}^{k'}$ defined by $f_{j_1}, \dots, f_{j_{k'}}$ is dominant, which exist by Lemma 4.3. Let $V_{f'} \subseteq \mathbb{A}^n_{\mathbb{F}(Y_1,\dots,Y_{k'})}$ be the generic fiber of f' (see the definition after Theorem 4.4). Finally, let $\ell_1, \dots, \ell_k \in \mathbb{F}[X_1, \dots, X_n]$ be linear polynomials such that both the morphism $\pi : V \to \mathbb{A}^k$ defined by $\ell_1, \dots, \ell_k$ and the morphism $\tau : V_{f'} \to \mathbb{A}^{k-k'}_{\mathbb{F}(Y_1,\dots,Y_{k'})}$ defined by $\ell_1, \dots, \ell_{k-k'}$ are finite.

Let $t \in \{0, \dots, k - k'\}$. Then there exists a polynomial $P \in \mathbb{F}[X_1, \dots, X_n]$ of degree at most $k' \cdot \deg V \cdot \prod_{i=1}^{k'} \deg h_i$ that does not vanish identically on V such that the following holds: Let $\varphi : \mathbb{A}^n \to \mathbb{A}^{t+m}$ be the polynomial map defined by $\ell_1, \dots, \ell_t, f_1, \dots, f_m$. Then for every $a \in V$ satisfying $P(a) \neq 0$, the fiber $\varphi|_V^{-1}(\varphi(a))$ is equidimensional of dimension $k - k' - t$.

As a corollary, we have the following effective fiber dimension theorem, stated in a more standard form.

Corollary 4.11 (Effective fiber dimension theorem – standard form)**.** *Let $V \subseteq \mathbb{A}^n$ be an irreducible affine variety over an algebraically closed field $\mathbb{F}$. Let $h_1, \dots, h_s \in \mathbb{F}[X_1, \dots, X_n]$ with $\deg h_1 \geq \cdots \geq \deg h_s$. Let $f_1, \dots, f_m \in \mathcal{L}_{h_1,\dots,h_s,\mathbb{F}}$, which define a polynomial map $f : \mathbb{A}^n \to \mathbb{A}^m$. Finally, let $W = \overline{f(V)} \subseteq \mathbb{A}^m$. Then there exists a polynomial $P \in \mathbb{F}[X_1, \dots, X_n]$ of degree at most $\dim W \cdot \deg V \cdot \prod_{i=1}^{\dim W} \deg h_i$ that does not vanish identically on V such that for every $a \in V$ satisfying $P(a) \neq 0$, the fiber $f|_V^{-1}(f(a))$ is equidimensional of dimension $\dim V - \dim W$.*

Proof. Use the notations in Theorem 4.10. Note that the linear polynomials $\ell_1, \ldots, \ell_k$ satisfying the conditions in Theorem 4.10 exist by Noether normalization (see Lemma 4.9). Now apply Theorem 4.10 with $t = 0$. □

Theorem 4.10 is proved in Appendix B.

Degree bound for the images of affine varieties. Finally, we need the following degree bound for the images of affine varieties (or more precisely, their closures) under polynomial maps.

Lemma 4.12. *Let $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ be an affine variety over a field $\mathbb{F}$. Let $h_1, \ldots, h_s \in \mathbb{F}[X_1, \ldots, X_n]$ with $\deg h_1 \geq \cdots \geq \deg h_s$. Let $f_1, \ldots, f_m \in \mathcal{L}_{h_1,\ldots,h_s,\mathbb{F}}$, which define a polynomial map $f : \mathbb{A}^n_{\mathbb{F}} \to \mathbb{A}^m_{\mathbb{F}}$. Finally, let $W = \overline{f(V)} \subseteq \mathbb{A}^m_{\mathbb{F}}$. Then*

$$\deg W \leq \deg V \cdot \prod_{i=1}^{\dim W} \deg h_i.$$

We prove Lemma 4.12 in Appendix C. It generalizes a bound in [BCS97, §8.5], which states that $\deg W \leq d^{\dim W}$ if $V = \mathbb{A}^n_{\mathbb{F}}$ and $\deg f_i \leq d \in \mathbb{N}^+$ for $i \in [m]$.

5 Linear Seeded Rank Extractors for Varieties

In this section, we consider the problem of constructing *seeded rank extractors for varieties* that are linear: i.e., a set of *linear* maps such that for every variety V most of the maps in the set preserve the dimension of V. We show that these objects are simply linear *seeded rank extractors for subspaces*, a well-known linear algebraic pseudorandom object for which explicit constructions were given in [GR08, FS12, For14].

The proof is based on the notion of *tangent spaces* of varieties, which are linear subspaces that are local first-order approximations of varieties. Intuitively, for an affine variety V, as we look at smaller and smaller neighborhoods of a *nonsingular point* a of V, the tangent space T_aV would become a better and better approximation of V. Thus, one should expect that a linear map that preserves the dimension of T_aV, which is a subspace, also preserves the dimension of V. While it is not entirely obvious what "smaller and smaller neighborhoods" mean in the Zariski topology, we will see that the claim is indeed true and follows from general facts in algebraic geometry.

Fix $\mathbb{F}$ to be an algebraically closed field throughout this section. We first formally define seeded rank extractors for varieties and subspaces.

Definition 5.1 (Seeded rank extractors). *Let $\varphi_1, \ldots, \varphi_\ell : \mathbb{A}^n \to \mathbb{A}^m$ be polynomial maps, where $n \geq m$. We say $(\varphi_i)_{i\in[\ell]}$ is an (n, m, k, ε)* seeded rank extractor for varieties *(resp. subspaces) if for every affine variety (resp. linear subspace) $V \subseteq \mathbb{A}^n$ over $\mathbb{F}$ of dimension at least k, all but at most ε-fraction of φ_i satisfy $\dim \overline{\varphi_i(V)} = m$ (or equivalently, $\varphi_i|_V : V \to \mathbb{A}^m$ is dominant). We call $\log \ell$ the* seed length *of the seeded rank extractor.*

In addition, we say $(\varphi_i)_{i\in[\ell]}$ is linear *if each φ_i is a linear map, i.e., defined by linear polynomials.*

The optimal choice of k is $k = m$, in which case the seeded rank extractor is "lossless." Explicit linear (n, m, k, ε) seeded rank extractors for subspaces with seed length $O(\log n + \log(1/\varepsilon))$ and $k = m$ was first constructed by Gabizon and Raz [GR08]. We use an improved construction given in [FS12, For14].

Lemma 5.2 ([FS12, For14])**.** *Let $n \in \mathbb{N}^+$ and $m \in [n]$. Let $\omega \in \mathbb{F}^\times$ such that the multiplicative order of ω is at least n. Let $s_1, \dots, s_\ell$ be distinct elements in $\mathbb{F}^\times$. For $i \in [\ell]$, let $\varphi_i : \mathbb{A}^n \to \mathbb{A}^m$ be the linear map defined by the $m \times n$ matrix $((\omega^{j'-1} s_i)^{j-1})_{j' \in [m], j \in [n]}$. In other words, φ_i is given by*

$$\varphi_i : (a_1, \dots, a_n) \mapsto \left(\sum_{j=1}^n s_i^{j-1} a_j, \sum_{j=1}^n (\omega s_i)^{j-1} a_j, \dots, \sum_{j=1}^n (\omega^{m-1} s_i)^{j-1} a_j \right).$$

Then $(\varphi_i)_{i \in [\ell]}$ is a linear (n, m, m, ε) seeded rank extractor for subspaces, where $\varepsilon = m(n-m)/\ell$.

The main result of this section is the following theorem.

Theorem 5.3. *An (n, m, k, ε) linear seeded rank extractor for subspaces is also an (n, m, k, ε) linear seeded rank extractor for varieties.*

Corollary 5.4. *The construction $(\varphi_i)_{i \in [\ell]}$ in Lemma 5.2 is a linear (n, m, m, ε) seeded rank extractor for varieties, where $\varepsilon = m(n-m)/\ell$.*

5.1 Tangent Spaces and the Jacobian Criterion for Smoothness

The proof of Theorem 5.3 uses the notion of *tangent spaces.*

Definition 5.5 (Tangent space)**.** *Let $V \subseteq \mathbb{A}^n$ be an affine variety over $\mathbb{F}$. For a point $a \in V$, the* tangent space *$T_a V$ of V at a is the linear subspace of $\mathbb{A}^n$ defined by the linear equations*

$$\sum_{i=1}^n \frac{\partial f}{\partial X_i}(a) \cdot X_i = 0, \qquad f = f(X_1, \dots, X_n) \in I(V).$$

To explain the intuition, note that at a point $a = (a_1, \dots, a_n) \in V$, the inhomogeneous linear equations $\sum_{i=1}^n \frac{\partial f}{\partial X_i}(a) \cdot (X_i - a_i) = 0$ with $f \in I(V)$ define an affine subspace of $\mathbb{A}^n$ passing through a, which can be seen as a first-order approximation of V locally at a. The tangent space $T_a V$ is defined to be the linear subspace that is a translate of this affine subspace.

The dimension of a tangent space is bounded from below by the dimension of the variety, as stated by the following lemma.

Lemma 5.6 ([Sha94, §II.1])**.** $\dim T_a V \geq \dim V$ *for every affine variety V over $\mathbb{F}$ and $a \in V$.*

The notion of *smoothness* can be defined in terms of whether the equality in Lemma 5.6 is attained. For simplicity, we define it only for irreducible affine varieties, which suffices for us.

Definition 5.7 (Smoothness)**.** *Let V be an irreducible affine variety over $\mathbb{F}$ and let $a \in V$. If $\dim T_a V = \dim V$, then we say V is* smooth *or* nonsingular *at a, and a is a* nonsingular point *of V. Otherwise, we say V is* non-smooth *or* singular *at a, and a is a* singular point *of V.*

Denote by V_{sing} the subset of singular points of V, called the singular locus *of V.*

Jacobian criterion for smoothness. The singular locus V_{sing} of V can be determined via the *Jacobian criterion*, which we explain now.

Definition 5.8 (Jacobian matrix). *Let $f_1, \ldots, f_m \in \mathbb{F}[X_1, \ldots, X_n]$. The associated* Jacobian matrix $J_{\mathbf{f}}$ *is an $m \times n$ matrix over the ring $\mathbb{F}[X_1, \ldots, X_n]$, defined by*

$$J_{\mathbf{f}} = \left(\frac{\partial f_i}{\partial X_j} \right)_{i \in [m], j \in [n]}.$$

Lemma 5.9. *Let $V \subseteq \mathbb{A}^n$ be an affine variety over $\mathbb{F}$, and let $f_1, \ldots, f_m \in \mathbb{F}[X_1, \ldots, X_n]$ such that $\{f_1, \ldots, f_m\}$ is a generating set of $I(V)$. Then for $a = (a_1, \ldots, a_n) \in V$, the tangent space $T_a V$ is the right nullspace of $J_{\mathbf{f}}(a)$.*

Proof. By definition, we want to show the following statement: $\sum_{j=1}^{n} \frac{\partial f_i}{\partial X_j}(a) \cdot a_j = 0$ for all $i \in [m]$ iff $\sum_{j=1}^{n} \frac{\partial f}{\partial X_j}(a) \cdot a_j = 0$ for all $f \in I(V)$. The "if" part is immediate as $f_1, \ldots, f_m \in I(V)$.

To see the "only if" part, consider $f \in I(V)$. Then $f = \sum_{i=1}^{m} g_i f_i$ for some $g_1, \ldots, g_m \in \mathbb{F}[X_1, \ldots, X_n]$. So for $j \in [n]$,

$$\frac{\partial f}{\partial X_j}(a) = \sum_{i=1}^{m} \left(\frac{\partial g_i}{\partial X_j}(a) f_i(a) + g_i(a) \frac{\partial f_i}{\partial X_j}(a) \right) = \sum_{i=1}^{m} g_i(a) \frac{\partial f_i}{\partial X_j}(a)$$

where the second equality holds as f_i vanishes at $a \in V$ for $i \in [m]$. It follows that if $\sum_{j=1}^{n} \frac{\partial f_i}{\partial X_j}(a) \cdot a_j = 0$ for $i \in [m]$, then $\sum_{j=1}^{n} \frac{\partial f}{\partial X_j}(a) \cdot a_j = 0$. □

Corollary 5.10 (Jacobian criterion for smoothness). *Let $V \subseteq \mathbb{A}^n$ be an irreducible affine variety of dimension k over $\mathbb{F}$, and let $f_1, \ldots, f_m \in \mathbb{F}[X_1, \ldots, X_n]$ such that $\{f_1, \ldots, f_m\}$ is a generating set of $I(V)$. Then the singular locus of V is given by*

$$V_{\mathrm{sing}} = \{a \in V : \operatorname{rank} J_{\mathbf{f}}(a) < n - k\},$$

which is an affine subvariety of V defined by the set of all $(n-k) \times (n-k)$ minors of $J_{\mathbf{f}}$.

Proof. By Lemma 5.9, the condition that $\operatorname{rank} J_{\mathbf{f}}(a) < n - k$ is equivalent to $\dim T_a V > k$. This is further equivalent to $\dim T_a V \neq k$ by Lemma 5.6. So the claim holds by definition. □

Remark. The above Jacobian criterion is related to but different from the *Jacobian criterion for algebraic independence*. The latter states that the transcendence degree of $\mathbb{F}(f_1, \ldots, f_m)/\mathbb{F}$ equals the rank of the associated Jacobian matrix $J_{\mathbf{f}}$, and in particular, $f_1, \ldots, f_m$ are algebraically independent iff $\operatorname{rank} J_{\mathbf{f}} = m$. However, this statement requires the characteristic of $\mathbb{F}$ to be zero or large, or more generally, a certain separability condition to hold. On the other hand, the Jacobian criterion for smoothness that we use holds without extra conditions. In particular, it works in *any* characteristic.

We also need the fact that varieties are almost-everywhere-nonsingular.

Lemma 5.11 ([Sha94, §II.1]). *The set of nonsingular points of an irreducible affine variety V over $\mathbb{F}$ is a dense open subset of V. That is, V_{sing} is a proper subvariety of V.*

5.2 Proof of Theorem 5.3

The proof of Theorem 5.3 is based on the following lemma.

Lemma 5.12. *Let $V \subseteq \mathbb{A}^n$ be an irreducible affine variety over $\mathbb{F}$ and let $a \in V$. Let $f_1, \ldots, f_m \in \mathbb{F}[X_1, \ldots, X_n]$, which defines a polynomial map $\varphi : \mathbb{A}^n \to \mathbb{A}^m$. Let W be the right nullspace of $J_{\mathbf{f}}(a) = \left(\frac{\partial f_i}{\partial X_j}(a)\right)_{i\in[m], j\in[n]}$. Suppose* $\dim(W \cap T_a V) = \dim V - m$. *Then* $\dim \overline{\varphi(V)} = m$.

Proof. Let $b = (b_1, \ldots, b_m) = \varphi(a) \in \mathbb{A}^m$, so that $a \in \varphi^{-1}(b)$. As $V \cap \varphi^{-1}(b)$ is a subvariety of both V and $\varphi^{-1}(b)$, we have $T_a(V \cap \varphi^{-1}(b)) \subseteq (T_a V) \cap (T_a \varphi^{-1}(b))$.

Note that the fiber $\varphi^{-1}(b) \subseteq \mathbb{A}^n$ is defined by the polynomials $\hat{f}_1, \ldots, \hat{f}_m$ where $\hat{f}_i := f_i - b_i$ for $i \in [m]$. So $\hat{f}_i \in I(\varphi^{-1}(b))$ for $i \in [m]$. Pick a finite generating set $\{g_1, \ldots, g_s\}$ of $I(\varphi^{-1}(b))$ such that $s \geq m$ and $g_i = \hat{f}_i$ for $i \in [m]$. By Lemma 5.9, the tangent space $T_a \varphi^{-1}(b)$ is the right nullspace of $J_{\mathbf{g}}(a) = \left(\frac{\partial g_i}{\partial X_j}(a)\right)_{i\in[s], j\in[n]}$. As $\frac{\partial \hat{f}_i}{\partial X_j} = \frac{\partial f_i}{\partial X_j}$ for $i \in [m]$ and $j \in [n]$, we have $J_{\mathbf{f}}(a) = \left(\frac{\partial \hat{f}_i}{\partial X_j}(a)\right)_{i\in[m], j\in[n]}$. As $g_i = \hat{f}_i$ for $i \in [m]$, the matrix $J_{\mathbf{f}}(a)$ is the upper $m \times n$ submatrix of $J_{\mathbf{g}}(a)$. It follows that $T_a \varphi^{-1}(b) \subseteq W$. So $T_a(V \cap \varphi^{-1}(b)) \subseteq W \cap T_a V$. Therefore,

$$\dim(V \cap \varphi^{-1}(b)) \leq \dim T_a(V \cap \varphi^{-1}(b)) \leq \dim(W \cap T_a V) = \dim V - m$$

where the first inequality holds by Lemma 5.6. On the other hand, by the fiber dimension theorem (Theorem 4.4),

$$\dim(V \cap \varphi^{-1}(b)) \geq \dim V - \dim \overline{\varphi(V)} \geq \dim V - m.$$

This forces $\dim(V \cap \varphi^{-1}(b)) = \dim V - m$ and $\dim \overline{\varphi(V)} = m$. □

We are now ready to prove Theorem 5.3.

Proof of Theorem 5.3. Let $V \subseteq \mathbb{A}^n$ be an affine variety over $\mathbb{F}$ of dimension at least k, and let V_0 be an irreducible component of V such that $\dim V_0 = \dim V$. Let a be a nonsingular point of V_0, which exists by Lemma 5.11. Then $\dim T_a V_0 = \dim V_0 \geq k$.

We claim that for a linear map $\varphi : \mathbb{A}^n \to \mathbb{A}^m$, if $\dim \varphi(T_a V_0) = m$, then $\dim \overline{\varphi(V_0)} = m$ and hence $\dim \overline{\varphi(V)} = m$. Note that this claim implies Theorem 5.3. This follows by choosing φ to be each of the linear maps in an (n, m, k, ε) linear seeded rank extractor and noting that $T_a V_0$ is a linear subspace of $\mathbb{A}^n$ of dimension at least k.

So it remains to prove the above claim. Assume $\dim \varphi(T_a V_0) = m$. Suppose φ is defined by linear polynomials $f_1, \ldots, f_m \in \mathbb{F}[X_1, \ldots, X_n]$. Let W be the kernel of φ. Then

$$\dim(W \cap T_a V_0) = \dim T_a V_0 - \dim \varphi(T_a V_0) = \dim V_0 - m.$$

As φ is linear, W equals the right nullspace of the matrix $J_{\mathbf{f}}(a) = \left(\frac{\partial f_i}{\partial X_j}(a)\right)_{i\in[m], j\in[n]}$. So $\dim \overline{\varphi(V_0)} = m$ by Lemma 5.12. This proves the claim and Theorem 5.3 follows. □

6 Deterministic Rank Extractors for Varieties

Let $\mathbb{F}$ be an algebraically closed field. In this section, we consider the problem of constructing explicit *deterministic (lossless) rank extractors/condensers for varieties.* These are polynomial maps $\mathbb{A}^n \to \mathbb{A}^m$ that preserve the dimension of low-degree affine varieties $V \subseteq \mathbb{A}^n$ over $\mathbb{F}$ but reduce the dimension of the ambient space.

Dvir, Gabizon and Wigderson [DGW09] constructed explicit deterministic rank extractors for *polynomial sources*. These objects can also be viewed as deterministic rank extractors for varieties that are the closures of the images of polynomial maps. A key technique used in their analysis is the *Jacobian criterion for algebraic independence*, which requires the characteristic of $\mathbb{F}$ to be zero or large.

To solve the problem for general varieties, one natural approach is generalizing the Jacobian criterion for algebraic independence. A key step in the proof of [DGW09] is showing that a certain polynomial associated with the Jacobian matrix is nonzero. Thus, it is natural for us to show that a similar polynomial does not vanish completely on affine varieties and that this is sufficient for constructing deterministic rank extractors for varieties.

While this idea can be made rigorous, the problem is that proving the nonvanishing of a polynomial on an affine variety appears to be challenging. We need to show that not only is the polynomial nonzero, but it remains nonzero modulo the ideal defining the variety. It is not clear to us how to prove such a result due to the generality of the variety.

The DKL construction. Instead of using a Jacobian-based construction, we take a different approach. Namely, we show that the explicit construction of *variety evasive sets* by Dvir, Kollár, and Lovett [DKL14] can be used to construct deterministic rank extractors for varieties. Variety evasive sets are large finite subsets of $\mathbb{A}^n$ that have small intersections with varieties of low degree and low dimension. While they do not give deterministic rank extractors for varieties in general, we show that the construction of variety evasive sets in [DKL14] does give such a construction.

More specifically, Dvir, Kollar and Lovett [DKL14] construct explicit variety evasive sets by constructing an explicit polynomial map $\varphi : \mathbb{A}^n \to \mathbb{A}^m$ defined by polynomials $f_1, \dots, f_m \in \mathbb{F}[X_1, \dots, X_n]$ such that the intersection of $\varphi^{-1}(\mathbf{0}) = V(f_1, \dots, f_m)$ with any low-degree variety of dimension at most m is finite, where $\mathbf{0}$ denotes the origin of $\mathbb{A}^n$. We observe that this remains true if $\varphi^{-1}(\mathbf{0})$ is replaced by $\varphi^{-1}(b)$ for any $b \in \mathbb{A}^m$. In other words, for any low-degree variety V of dimension at most m, the polynomial map φ restricts to a morphism $\varphi|_V : V \to \mathbb{A}^m$ whose fibers are all finite sets. In the terminology of algebraic geometry, this means $\varphi|_V$ is a *quasi-finite morphism*. By the fiber dimension theorem (Theorem 4.4), we then have $\dim \overline{\varphi(V)} = \dim(V)$.

In this section, we construct explicit deterministic rank extractors and rank condensers for varieties by adapting the analysis in [DKL14]. We also formulate the construction in a way that highlights the connection with linear error-correcting codes. In particular, a linear MDS code yields a deterministic rank extractor for varieties in the sense that the coefficients of the polynomials that define the rank extractor are specified by a parity-check matrix of the code.

In Section 11 and Appendix D, we will show that the polynomial map φ has the stronger property that $\varphi|_V$ is a *finite morphism*, not just quasi-finite, and this gives explicit Noether normalization lemmas for affine varieties and affine algebras.

6.1 The Explicit Construction

We first define deterministic rank extractors and rank condensers for varieties.

Definition 6.1 (Deterministic rank extractors/condensers for varieties). *Let $n \in \mathbb{N}^+$ and $m \in [n]$. A polynomial map $\varphi : \mathbb{A}^n \to \mathbb{A}^m$ is an (n, m, k, d)* deterministic (lossless) rank condenser *if*

$\dim \overline{\varphi(V)} = \dim V$ *for every affine variety $V \subseteq \mathbb{A}^n$ over $\mathbb{F}$ of dimension at most k and degree at most d. When $k = m$, we also say φ is an (n, m, d)* deterministic (lossless) rank extractor*.*

k-regular matrices. Let $n \in \mathbb{N}^+$ and $m, k \in [n]$. We say a matrix $M \in \mathbb{F}^{m\times n}$ is *k-regular* if any k distinct columns of M are linearly independent. (The same definition was given in [DKL14] but for only for the special case where $k = m$.)

The following lemma gives a coding-theoretic characterization of k-regularity. Its proof is straightforward.

Lemma 6.2. *Let $\mathbb{K}$ be a subfield of $\mathbb{F}$ and let $M \in \mathbb{K}^{m\times n} \subseteq \mathbb{F}^{m\times n}$, where $n \in \mathbb{N}^+$ and $m, k \in [n]$. The following statements hold.*

- *M is k-regular iff there does not exist a nonzero vector $u \in \mathbb{K}^n$ of Hamming weight at most k such that $Mu = 0$.*
- *Suppose $k = m$. Then M is k-regular iff it is an* MDS matrix*, i.e., every maximal minor of M is nonzero.*

In particular, assuming $\mathbb{K}$ is a finite field, the matrix M is k-regular iff the linear code $C = \{u \in \mathbb{K}^n : Mu = 0\}$ over $\mathbb{K}$ defined by the parity check matrix M has minimum distance at least $k+1$. And if $k = m$, then M is k-regular iff C is a linear MDS code of minimum distance $k+1$, i.e., it is a linear code of dimension $n-k$ and minimum distance $k+1$.[3]

The construction. We now present the explicit construction of deterministic rank extractors and condensers for varieties. It is based on the explicit construction of variety evasive sets in [DKL14].

Let $n, d \in \mathbb{N}^+$ and $m, k \in [n]$. Let $d_1, \ldots, d_n$ be n pairwise coprime integers greater than d.[4] Let $M = (c_{i,j})_{i\in[m],j\in[n]} \in \mathbb{F}^{m\times n}$ be a k-regular matrix. Let $\varphi = \varphi(M) : \mathbb{A}^n \to \mathbb{A}^m$ be the polynomial map

$$\varphi : (a_1, \ldots, a_n) \mapsto \left(\sum_{j=1}^{n} c_{1,j} a_j^{d_j}, \ldots, \sum_{j=1}^{n} c_{m,j} a_j^{d_j} \right).$$

We remark that, curiously, the construction above is very similar to the construction of an affine extractor in Section 10, although their purposes and the techniques used to analyze them are substantially different.

The following theorem and its corollaries are the main results of this section.

Theorem 6.3. *For every $b \in \mathbb{A}^m$ and every affine variety $V \subseteq \mathbb{A}^n$ over $\mathbb{F}$ of dimension at most k and degree at most d, the fiber $(\varphi|_V)^{-1}(b) = \varphi^{-1}(b) \cap V$ is a finite set.*

Corollary 6.4. *φ is an (n, m, k, d) deterministic rank condenser for varieties. In particular, if $m = k$, then φ is an (n, m, d) deterministic rank extractor for varieties.*

Proof (assuming Theorem 6.3). Let $V \subseteq \mathbb{A}^n$ be an affine variety over $\mathbb{F}$ of dimension at most k and degree at most d. Let V_0 be an irreducible component of V. Then $\dim V_0 \le k$ and $\deg V_0 \le d$. It suffices to show $\dim \overline{\varphi(V_0)} = \dim V_0$. This follows from Theorem 6.3 and the fiber dimension theorem (Theorem 4.4). □

[3] We define the minimum distance of the zero code $\{0\}$ to be $n+1$, so that the statement also holds for $k = n$.

[4] While [DKL14] assumes $d_1 > \cdots > d_n$, this assumption does not really matter.

Choosing $d_1, \dots, d_n$. We still need to argue that $d_1, \dots, d_n$ and M can be computed efficiently. One can choose $d_1, \dots, d_n$ to be n distinct primes greater than d. The resulting deterministic time complexity of computing these integers is $\operatorname{poly}(n, d)$. The polynomial dependence on d is due to the fact that there is no known deterministic $N^{o(1)}$-time algorithm for finding primes greater than an integer $N > 0$.

To improve the time complexity, we may compute d_i in the following alternative way. Compute the smallest n distinct primes $p_1, \dots, p_n$, which have order $O(n \log n)$. For $i \in [n]$, let d_i be the smallest power of p_i such that $d_i > d$, so that $d_i = O(p_i d) = O(nd \log n)$. Then $d_1, \dots, d_n$ can be computed in time $\operatorname{poly}(n, \log d)$.

Choosing the matrix M. We need to choose a k-regular matrix M. For the problem of constructing an (n, m, d) deterministic rank extractor for varieties (i.e., $k = m$), we need to choose $M \in \mathbb{F}^{n \times m}$ to be an MDS matrix by Lemma 6.2. This can be achieved by choosing M to be an Vandermonde matrix $(\omega_j^{i-1})_{i \in [m], j \in [n]}$ with distinct $\omega_1, \dots, \omega_n \in \mathbb{F}$.

Suppose $\mathbb{F}$ has a finite subfield $\mathbb{F}_q$. Then using a Vandermonde matrix, we need $q \geq n$ to have $M \in \mathbb{F}_q^{m \times n}$. The condition $q \geq n$ can be relaxed to $q \geq n - 1$ in general, and to $q \geq n - 2$ in some special cases as explicit MDS matrices $M \in \mathbb{F}_q^{m \times n}$ are known in these cases [MS77].

In the case where $k = m \in \{1, n-1, n\}$, we do not need any lower bound on q as all-one vectors and identity matrices are always MDS matrices, and so is the $(n-1) \times n$ matrix $M = (c_{i,j})_{i \in [n-1], j \in [n]}$ defined by

$$c_{i,j} = \begin{cases} 1 & i = j, \\ -1 & j = n, \\ 0 & \text{otherwise.} \end{cases}$$

So we have the following corollary.

Corollary 6.5. *For $m \in \{1, n-1, n\}$, there exists an explicit construction of an (n, m, d) deterministic rank extractor for varieties that is defined by polynomials $f_1, \dots, f_m \in \mathbb{F}[X_1, \dots, X_n]$ satisfying the following:*

- *All the coefficients of $f_1, \dots, f_m$ are in $\{0, 1, -1\}$, and hence are in every subfield of $\mathbb{F}$.*
- $\deg f_1, \dots, \deg f_m = O((n+d)\log(n+d))$*. And the sparse representations of $f_1, \dots, f_m$ can be computed in time* $\operatorname{poly}(n, d)$*. The time complexity can be improved to* $\operatorname{poly}(n, \log d)$ *at the cost of increasing the degrees of $f_1, \dots, f_m$ to $O(nd \log n)$.*

A similar statement holds for general $m \in [n]$ and the coefficients of $f_1, \dots, f_m$ can be chosen in a finite field $\mathbb{F}_q$, assuming $\mathbb{F}_q$ is a subfield of $\mathbb{F}$ and $q \geq n - 1$. The time complexity would also depend polynomially on $\log q$.

The above explicit (n, m, d) deterministic extractor for varieties will be used in the proof of Theorem 1, but only in the case where $m = 1$. Previously, Dvir [Dvi12, Theorem 3.1] gave an explicit construction of an $(n, 1, d)$ deterministic rank extractor for varieties, where the polynomial defining the rank extractor is recursively constructed and has degree $\operatorname{poly}(d^n)$. Corollary 6.5 improves the degree of the polynomial to $\widetilde{O}(n + d)$ or $\widetilde{O}(nd)$.

6.2 Proof of Theorem 6.3

Theorem 6.3 can be proved via a simple adaptation of the proof in [DKL14]. For the sake of completeness, we present the proof below.

First, we need the following two lemmas from [DKL14].

Lemma 6.6 ([DKL14], Lemma 3.1)**.** *Let $V \subseteq \mathbb{A}^n$ be an affine variety over $\mathbb{F}$ of dimension at least one. Then there exist Laurent power series $h_1(T), \ldots, h_n(T) \in \mathbb{F}((T))$ such that*

1. *at least one $h_i(T)$ has a pole, i.e., $h_i(T) \notin \mathbb{F}[[T]]$, and*
2. *$P(h_1(T), \ldots, h_n(T)) = 0$ for all $P \in I(V)$.*

Lemma 6.7 ([DKL14], Lemma 3.5)**.** *Let $V \subseteq \mathbb{A}^n$ be an affine variety over $\mathbb{F}$ of dimension $k < n$ and degree d. For every $J = \{i_1, \ldots, i_{k+1}\} \subseteq [n]$ of size $k+1$, there exists a nonzero polynomial $g \in I(V) \cap \mathbb{F}[X_{i_1}, \ldots, X_{i_{k+1}}]$ of degree at most d.*

We adapt the proof of [DKL14] to prove the following lemma.

Lemma 6.8. *Let $M = (c_{i,j})_{i\in[m],j\in[n]} \in \mathbb{F}^{m\times n}$ be a k-regular matrix, and let $d_1, \ldots, d_n$ be pairwise coprime integers greater than d. Let $V \subseteq \mathbb{A}^n$ be an affine variety over $\mathbb{F}$ of dimension at most k and degree at most d. Then, there do not exist $h_1(T), \ldots, h_n(T) \in \mathbb{F}((T))$ that simultaneously satisfy the following conditions:*

1. *At least one $h_i(T)$ has a pole, i.e., $h_i(T) \notin \mathbb{F}[[T]]$.*
2. *$P(h_1(T), \ldots, h_n(T)) = 0$ for all $P \in I(V)$.*
3. *$\sum_{j=1}^{n} c_{i,j} h_j(T)^{d_j} \in \mathbb{F}[[T]]$ for all $i \in [m]$.*

Proof. By replacing k with $k' = \dim V$, we may assume the dimension of V is exactly k. Assume to the contrary that there exist $h_1(T), \ldots, h_n(T) \in \mathbb{F}((T))$ satisfying the three conditions. Let R be the greatest integer such that the term T^{-R} appears in the Laurent series $h_j(T)^{d_j}$ for some $j \in [n]$. By the first condition, at least one of the $h_j(T)$ has a pole, so we know $R > 0$.

Let J be the set of $j \in [n]$ for which the term T^{-R} appears in the Laurent series $h_j(T)^{d_j}$. Then $J \neq \emptyset$. We claim $|J| \geq k+1$. To see this, let $u_j \in \mathbb{F}$ be the coefficient of the term T^{-R} in $h_j(T)^{d_j}$ for $j \in [n]$. Let $u = (u_1, \ldots, u_n) \in \mathbb{F}^n$. Then the support of u is precisely J. By the third condition, for $i \in [m]$,

$$\sum_{j=1}^{n} c_{i,j} h_j(T)^{d_j} \in \mathbb{F}[[T]]. \tag{2}$$

The coefficient of T^{-R} in the LHS of (2) is $\sum_{j=1}^{n} c_{i,j} u_j$, which equals zero by (2) and the fact that $R > 0$. So we get the equation

$$Mu = 0.$$

By the k-regularity of M and Lemma 6.2, the Hamming weight of u is at least $k+1$, i.e., $|J| \geq k+1$. This proves the claim. Also note that this implies $k < n$ as $|J| \leq n$.

From here, the rest of the proof is identical to that of [DKL14, Theorem 2.1]. We present the proof for the sake of completeness. For $j \in J$, let r_j be the maximal integer such that T^{-r_j} appears in h_j, and it follows that $r_j = R/d_j$.

Let $\{j_1, \ldots, j_{k+1}\}$ be a subset of J of size $k+1$. By Lemma 6.7, there exists a nonzero polynomial $g(X_{j_1}, \ldots, X_{j_{k+1}}) \in I(V) \cap \mathbb{F}[X_{j_1}, \ldots, X_{j_{k+1}}]$ of degree at most d. By the second condition, we have

$$g(h_{j_1}(T), \ldots, h_{j_{k+1}}(T)) = 0. \tag{3}$$

Next, we observe that for every monomial $Q = X_{j_1}^{\gamma_1} \cdots X_{j_{k+1}}^{\gamma_{k+1}}$, the term $T^{-\sum_{i=1}^{k+1} \gamma_i r_{j_i}}$ appears in $Q(h_{j_1}(T), \ldots, h_{j_{k+1}}(T))$. Choose a monomial $X_{j_1}^{\alpha_1} \cdots X_{j_{k+1}}^{\alpha_{k+1}}$ that appears in g such that $\sum_{i=1}^{k+1} \alpha_i r_{j_i}$ is maximized. Such a monomial exists as $g \neq 0$. Then g must contain a different monomial $X_{j_1}^{\beta_1} \cdots X_{j_{k+1}}^{\beta_{k+1}}$ such that

$$\sum_{i=1}^{k+1} \alpha_i r_{j_i} = \sum_{i=1}^{k+1} \beta_i r_{j_i}. \tag{4}$$

Otherwise, the term $T^{-\sum_{i=1}^{k+1} \alpha_i r_{j_i}}$ would appear in $g(h_{j_1}(T), \ldots, h_{j_{k+1}}(T))$, which contradicts (3).

Let $D = \prod_{i=1}^{k+1} d_{j_i}$. Plugging $r_{j_i} = R/d_{j_i}$ into (4) and then multiplying both sides of (4) by D/R, we get

$$\sum_{i=1}^{k+1} \alpha_i D/d_{j_i} = \sum_{i=1}^{k+1} \beta_i D/d_{j_i}. \tag{5}$$

Consider arbitrary $i \in [k+1]$. Taking (5) modulo d_{j_i}, we get

$$\alpha_i D/d_{j_i} \equiv \beta_i D/d_{j_i} \pmod{d_{j_i}}.$$

As D/d_{j_i} is coprime to d_{j_i}, we may cancel it from both sides, which gives

$$\alpha_i \equiv \beta_i \pmod{d_{j_i}}.$$

As $0 \leq \alpha_i, \beta_i \leq \deg(g) \leq d < d_{j_i}$, we have $\alpha_i = \beta_i$. As $i \in [k+1]$ is arbitrary, we have $(\alpha_1, \ldots, \alpha_{k+1}) = (\beta_1, \ldots, \beta_{k+1})$, contradicting the assumption that $X_{j_1}^{\alpha_1} \cdots X_{j_{k+1}}^{\alpha_{k+1}} \neq X_{j_1}^{\beta_1} \cdots X_{j_{k+1}}^{\beta_{k+1}}$. □

Remark. Lemma 6.8 was implicitly proved in [DKL14] except that the third condition was replaced by the stronger statement $\sum_{j=1}^{n} c_{i,j} h_j(T)^{d_j} = 0$. Our observation is that the proof still works if this condition is relaxed to $\sum_{j=1}^{n} c_{i,j} h_j(T)^{d_j} \in \mathbb{F}$, or even $\sum_{j=1}^{n} c_{i,j} h_j(T)^{d_j} \in \mathbb{F}[[T]]$. (As can be seen below, the former relaxation suffices for proving Theorem 6.3, but we will need the latter in Section 11 and Appendix D when we prove that $\varphi|_V$ is a finite morphism.)

Now we are ready to prove Theorem 6.3.

Proof of Theorem 6.3. Assume to the contrary that Theorem 6.3 does not hold. Then there exist an affine variety $V \subseteq \mathbb{A}^n$ over $\mathbb{F}$ of dimension at most k and degree at most d and $b = (b_1, \ldots, b_m) \in \mathbb{A}^m$ such that $\varphi^{-1}(b) \cap V$ is not finite, i.e., $\dim(\varphi^{-1}(b) \cap V) \geq 1$. Applying Lemma 6.6 to $\varphi^{-1}(b) \cap V$, we see that there exist Laurent power series $h_1(T), \ldots, h_n(T) \in \mathbb{F}((T))$ such that

1. at least one $h_i(T)$ has a pole, and
2. $P(h_1(T), \ldots, h_n(T)) = 0$ for all $P \in I(\varphi^{-1}(b) \cap V)$.

As $I(V) \subseteq I(\varphi^{-1}(b) \cap V)$, the second item implies $P(h_1(T), \ldots, h_n(T)) = 0$ for all $P \in I(V)$. In addition, for $i \in [m]$, we have $\left(\sum_{j=1}^n c_{i,j} X_j^{d_j}\right) - b_i \in I(\varphi^{-1}(b)) \subseteq I(\varphi^{-1}(b) \cap V)$ by the definition of φ. So the second item also implies

$$\sum_{j=1}^n c_{i,j} h_j(T)^{d_j} = b_i \in \mathbb{F} \subseteq \mathbb{F}[[T]] \qquad \text{for } i \in [m].$$

But then $h_1(T), \ldots, h_n(T)$ satisfy the three conditions in Lemma 6.8, and $M = (c_{i,j})_{i \in [m], j \in [n]}$ is k-regular. This contradicts Lemma 6.8. □

7 Decomposition and Min-Entropy Estimation of (n, k, d) Algebraic Sources

In this section, we prove that every (n, k, d) algebraic source can be (approximately) decomposed into a convex combination of irreducible, or even irreducibly minimal (n, k, d) sources. In particular, this reduces the problem of constructing deterministic extractors for general (n, k, d) algebraic sources to that for irreducibly minimal (n, k, d) algebraic sources. We will use this reduction in Section 8.

In addition, we show that every (n, k, d) algebraic source D over $\mathbb{F}_q$ is close to a distribution with min-entropy about $k \log q$, and that this estimation is tight up to an additive term of order $O(\log d)$ assuming that k is maximized, i.e., that D is not an $(n, k+1, d)$ algebraic source over $\mathbb{F}_q$.

7.1 Decomposition of (n, k, d) Algebraic Sources

First, we prove some useful lemmas.

Lemma 7.1. *Let $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ be an affine variety of dimension k over a field $\mathbb{F}$. Let $V_1, \ldots, V_s$ be the irreducible components of V. Suppose S is a finite subset of $\mathbb{F}$. For $i \in [s]$, let B_i be the subset of $V_i \cap S^n$ consisting of the points that are in the intersection of at least two irreducible components of V. Then $\sum_{i=1}^s |B_i| \leq (\deg V)^2 |S|^{k-1}$.*

Proof. For $i \in [s]$, we have $|B_i| \leq \deg V_i \cdot \deg V \cdot q^{k-1}$ by Bézout's inequality (Lemma 4.2) and Lemma 4.5. So $\sum_{i=1}^s |B_i| \leq \sum_{i=1}^s \deg V_i \cdot \deg V \cdot q^{k-1} = (\deg V)^2 \cdot q^{k-1}$. □

Lemma 7.2. *Let $V \subseteq \mathbb{A}^n_{\mathbb{F}_q}$ be an affine variety of dimension k over $\mathbb{F}_q$ such that no irreducible component of V is absolutely irreducible. Then $|V(\mathbb{F}_q)| \leq (\deg V)^2 q^{k-1}$.*

Proof. Naturally identify $V(\mathbb{F}_q)$ with $V_{\overline{\mathbb{F}}_q} \cap \mathbb{F}_q^n$. The Frobenius automorphism over $\mathbb{F}_q$ sends every irreducible component of $V_{\overline{\mathbb{F}}_q}$ to a different irreducible component as these irreducible components are not absolutely irreducible. But it fixes every point in $V(\mathbb{F}_q)$ since these points are rational. So every point in $V(\mathbb{F}_q)$ is in the intersection of at least two irreducible components of $V_{\overline{\mathbb{F}}_q}$. Applying Lemma 7.1 with $\mathbb{F} = \overline{\mathbb{F}}_q$ and $S = \mathbb{F}_q$, we see $|V(\mathbb{F}_q)| \leq (\deg V_{\overline{\mathbb{F}}_q})^2 q^{k-1} = (\deg V)^2 q^{k-1}$. □

Lemma 7.3. *Suppose S is a finite set and B is a proper subset of S. Let D and D' be the uniform distributions over S and $S \setminus B$ respectively. Then D and D' are $\frac{|B|}{|S|}$-close.*

Proof. Consider the following process: First sample $x \sim D$. If $x \notin B$, then output x. Otherwise sample $y \sim D'$ and output y. The distribution of the output is exactly D'. So the statistical distance between D and D' is at most $\Pr_{x \sim D}[x \in B] \le \frac{|B|}{|S|}$. □

The next lemma states that every (n, k, d) algebraic source D can be approximately decomposed into a convex combination of irreducible (n, k, d) algebraic sources, and such a decomposition preserves the minimality (i.e., the property that $\dim V = k$). The idea behind the proof is quite natural: we start by decomposing V as a union of irreducible components. We then observe that the contribution of components that are not absolutely irreducible or have dimensions strictly lower than $\dim V$ is small (using Lemma 7.2). In addition, the remaining irreducible components are approximately disjoint, namely, the intersection of any two distinct irreducible components is small (by Lemma 7.1). This implies that these irreducible components approximately define a convex combination of irreducible algebraic sources that is close to D.

Lemma 7.4 (Decomposition into irreducible algebraic sources)**.** *Suppose $q \ge \max\{20d^5, 2d^2/\varepsilon\}$, where $\varepsilon \in (0, 1)$. Then every (n, k, d) algebraic source D over $\mathbb{F}_q$ is ε-close to a convex combination of irreducible (n, k, d) algebraic sources D_i over $\mathbb{F}_q$. Moreover, if D is a minimal (n, k, d) algebraic source over $\mathbb{F}_q$, then each D_i can be chosen to be an irreducibly minimal (n, k, d) algebraic source over $\mathbb{F}_q$.*

Proof. Let D be an (n, k, d) algebraic source. Let $V \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ and $f : \mathbb{A}^r_{\mathbb{F}_q} \to \mathbb{A}^n_{\mathbb{F}_q}$ be as in Definition 1.2 so that $D = f(U_{V(\mathbb{F}_q)})$.

Let V_* be the union of the irreducible components of V that are absolutely irreducible and have dimension $\dim V$, and let V_*^{c} be the union of the remaining irreducible components of V. Then $V_* \neq \emptyset$ by the first condition in Definition 1.2. By the effective Lang–Weil bound (Theorem 4.6) and the assumption that $q \ge \max\{20d^5, 2d^2/\varepsilon\}$, we have

$$|V(\mathbb{F}_q)| \ge |V_*(\mathbb{F}_q)| \ge q^{\dim V}/2 \ge d^2 q^{\dim V - 1}/\varepsilon. \tag{6}$$

Consider an irreducible component V_0 of V_*^{c}. Either $\dim V_0 < \dim V$ holds or V_0 is not absolutely irreducible. In the former case, we have $|V_0(\mathbb{F}_q)| \le \deg V_0 \cdot q^{\dim V - 1}$ by Lemma 4.5. And in the latter case, we have $|V_0(\mathbb{F}_q)| \le (\deg V_0)^2 \cdot q^{\dim V - 1}$ by Lemma 7.2. Summing over all V_0, we conclude that

$$|V_*^{\mathrm{c}}(\mathbb{F}_q)| \le (\deg V_*^{\mathrm{c}})^2 \cdot q^{\dim V - 1}. \tag{7}$$

Let $U_{V_*(\mathbb{F}_q)}$ be the uniform distribution over $V_*(\mathbb{F}_q)$. By Lemma 7.3, the distributions $U_{V(\mathbb{F}_q)}$ and $U_{V_*(\mathbb{F}_q)}$ are ε'-close, where

$$\varepsilon' := \frac{|V(\mathbb{F}_q) \setminus V_*(\mathbb{F}_q)|}{|V(\mathbb{F}_q)|} \le \frac{|V_*^{\mathrm{c}}(\mathbb{F}_q)|}{|V(\mathbb{F}_q)|} \overset{(6),(7)}{\le} \frac{(\deg V_*^{\mathrm{c}})^2}{d^2} \cdot \varepsilon. \tag{8}$$

Let $V_1, \dots, V_s$ be the irreducible components of V_*. For $i \in [s]$, let U_i be the uniform distribution over $V_i(\mathbb{F}_q)$. Let $N = \sum_{i=1}^s |V_i(\mathbb{F}_q)| \ge |V_*(\mathbb{F}_q)|$. Define the distribution U' to be the convex combination

$$U' := \sum_{i=1}^s \frac{|V_i(\mathbb{F}_q)|}{N} U_i.$$

Let B be the set of points in $V_*(\mathbb{F}_q)$ that are in the intersection of at least two irreducible components of V_*, and let $B_i = B \cap V_i(\mathbb{F}_q)$ for $i \in [s]$. Note that all the points in $V_*(\mathbb{F}_q) \setminus B$ have the same probability $\frac{1}{N}$ in the distribution U'.

Let $p_B = \Pr_{x\sim U'}[x \in B] = \frac{\sum_{i=1}^{s}|B_i|}{N}$. Consider the following process: Sample $x \sim U'$. If $x \notin B$, then output x. Otherwise, output $x' \in V_*(\mathbb{F}_q)$ such that each $y \in V_*(\mathbb{F}_q)$ is output with probability p_y, where

$$p_y = \begin{cases} p_B^{-1} \Pr_{x\sim U_{V_*(\mathbb{F}_q)}}[x = y] & \text{if } y \in B, \\ p_B^{-1}\left(\Pr_{x\sim U_{V_*(\mathbb{F}_q)}}[x = y] - \frac{1}{N}\right) = p_B^{-1}\left(\frac{1}{|V_*(\mathbb{F}_q)|} - \frac{1}{N}\right) & \text{if } y \notin B. \end{cases}$$

It is easy to verify that the probabilities p_y do define a distribution over $V_*(\mathbb{F}_q)$. Moreover, they are chosen in the way that the output distribution of the above process is precisely $U_{V_*(\mathbb{F}_q)}$. It follows that $U_{V_*(\mathbb{F}_q)}$ and U' are p_B-close. By Lemma 7.1, we have $\sum_{i=1}^{s}|B_i| \le (\deg V_*)^2 q^{\dim V - 1}$. So

$$p_B = \frac{\sum_{i=1}^{s}|B_i|}{N} \le \frac{(\deg V_*)^2 q^{\dim V-1}}{N} \le \frac{(\deg V_*)^2 q^{\dim V-1}}{|V_*(\mathbb{F}_q)|} \overset{(6)}{\le} \frac{(\deg V_*)^2}{d^2} \cdot \varepsilon. \tag{9}$$

As $\deg V_* + \deg V_*^{\mathrm{c}} = \deg V \le d$, we have

$$\varepsilon' + p_B \overset{(8),(9)}{\le} \frac{(\deg V_*^{\mathrm{c}})^2}{d^2} \cdot \varepsilon + \frac{(\deg V_*)^2}{d^2} \cdot \varepsilon \le \varepsilon.$$

So $U_{V(\mathbb{F}_q)}$ and U' are ε-close. It follows that $D = f(U_{V(\mathbb{F}_q)})$ and $f(U')$ are ε-close. Recall that U' is a convex combination of $U_1, \dots, U_s$, where each U_i is the uniform distribution over $V_i(\mathbb{F}_q)$. And by definition, $f(U_i)$ is an irreducible (n, k, d) algebraic sources over $\mathbb{F}_q$ for $i \in [s]$. It follows that D is ε-close to a convex combination of the irreducible (n, k, d) algebraic sources $f(U_1), \dots, f(U_s)$ over $\mathbb{F}_q$.

Finally, if D is a minimal (n, k, d) algebraic source over $\mathbb{F}_q$, then by definition, the affine variety V may be chosen such that $\dim V = k$. Then we also have $\dim V_i = k$ for $i \in [s]$ in the above proof. In this case, each $f(U_i)$ is an irreducibly minimal (n, k, d) algebraic source over $\mathbb{F}_q$ by definition. So D is ε-close to a convex combination of irreducibly minimal (n, k, d) algebraic sources over $\mathbb{F}_q$. □

Next, we further decompose an irreducible (n, k, d) algebraic source into a convex combination of irreducibly minimal (n, k, d) algebraic sources. Our main tool is the effective fiber dimension theorem (Theorem 4.10). Using this theorem and the results of Section 4, we intersect the variety V with various translates of a carefully chosen linear subspace. There are some bad events that could happen for some of these intersections. For example, the intersection may have the "wrong" dimension, or the resulting variety might have the "correct" dimension k but none of the irreducible components of dimension k are absolutely irreducible. Using the effective fiber dimension theorem, we are able to show that these bad events correspond to small portions of the variety V, and then we again obtain a natural way to decompose the remaining part as a convex combination of irreducibly minimal (n, k, d) sources.

We start with the following application of the fiber dimension theorem.

Lemma 7.5. *Let $V \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ be an irreducible affine variety over $\mathbb{F}_q$. Let $\varphi = (\varphi_1, \varphi_2) : V \to \mathbb{A}^n_{\mathbb{F}_q}$ be a dominant morphism defined by $f_1, \dots, f_n$, where φ_1 and φ_2 are defined by the first n_1 and the last n_2 polynomials respectively and $n_1 + n_2 = n$. Let U be the subset of $a \in V(\mathbb{F}_q)$ such that $\varphi^{-1}(\varphi(a))$ is equidimensional of dimension $\dim V - n_1 - n_2$. Then we have:*

1. *For $a \in U$, the fiber $V_{\varphi_1(a)} := \varphi_1^{-1}(\varphi_1(a))$ is equidimensional of dimension* $\dim V - n_1$.

2. *For $a \in U$ and each irreducible component Z of $V_{\varphi_1(a)}$, we have* $\dim \overline{\varphi_2(Z)} \leq n_2$. *Moreover, the equality is attained if $a \in Z$.*

Proof. Let $a \in U$. By the fiber dimension theorem (Theorem 4.4), every irreducible component of $V_{\varphi_1(a)}$ is at least $\dim V - \dim \overline{\varphi_1(V)} \geq \dim V - n_1$. Also note that $\varphi^{-1}(\varphi(a)) = \varphi_2|_{V_{\varphi_1(a)}}^{-1}(\varphi_2(a))$. Again by the fiber dimension theorem, we have

$$\dim \varphi^{-1}(\varphi(a)) \geq \dim V_{\varphi_1(a)} - \dim \overline{\varphi_2(V_{\varphi_1(a)})} \geq \dim V_{\varphi_1(a)} - n_2.$$

We know $\dim \varphi^{-1}(\varphi(a)) = \dim V - n_1 - n_2$ by assumption. So $\dim V_{\varphi_1(a)} \leq (\dim V - n_1 - n_2) + n_2 = \dim V - n_1$. It follows that $V_{\varphi_1(a)}$ is equidimensional of dimension $\dim V - n_1$, which proves the first claim.

Let Z be an irreducible component of $V_{\varphi_1(a)}$. We already know that

$$\dim \overline{\varphi_2(Z)} \leq \dim \overline{\varphi_2(V_{\varphi_1(a)})} \leq n_2.$$

Now assume $a \in Z$. Let W be the irreducible component of $\varphi^{-1}(\varphi(a))$ that contains a. We have $\dim W = \dim V - n_1 - n_2$ by assumption and $\dim Z = \dim V - n_1$ by the first claim. Note that W is an irreducible component of $\varphi_2|_Z^{-1}(\varphi_2(a))$. So by the fiber dimension theorem,

$$\dim V - n_1 - n_2 = \dim W \geq \dim Z - \dim \overline{\varphi_2(Z)} = (\dim V - n_1) - \dim \overline{\varphi_2(Z)}$$

which implies that $\dim \overline{\varphi_2(Z)} \geq n_2$. So $\dim \overline{\varphi_2(Z)} = n_2$. □

Consider the setup in the previous lemma and further assume that V is absolutely irreducible. The following lemma roughly asserts that for most values of b in the image of φ_1, the fiber V_b satisfies the property that all but at most an ε fraction of its points come from irreducible components Z such that $\dim Z = \dim V - n_1$, Z is absolutely irreducible, and $\dim \overline{\varphi_2(Z)} = n_2$. In other words, the set of "bad" points in V_b that belong to other irreducible components is negligible.

Lemma 7.6. *Let $V \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ be an absolutely irreducible affine variety over $\mathbb{F}_q$. Let $\varphi = (\varphi_1, \varphi_2) : V \to \mathbb{A}^n_{\mathbb{F}_q}$ be a dominant morphism defined by $f_1, \dots, f_n$, where φ_1 and φ_2 are defined by the first n_1 and the last n_2 polynomials respectively and $n_1 + n_2 = n$. For $b \in \mathbb{F}_q^{n_1}$, let $V_b = \varphi_1^{-1}(b)$, and let V'_b be the union of the irreducible components Z of V_b such that Z is absolutely irreducible of dimension* $\dim V - n_1$ *and* $\dim \overline{\varphi_2(Z)} = n_2$. *Define*

$$\delta = \Pr_{a \sim U_{V(\mathbb{F}_q)}} [\dim \varphi^{-1}(\varphi(a)) \neq \dim V - n].$$

Let $d \in \mathbb{N}^+$ and $\varepsilon = (2d^2/q + \delta)^{1/2}$. Assume $q \geq 20d^5$, $\deg V \leq d$, *and* $\deg V_b \leq d$ *for all $b \in \varphi_1(V(\mathbb{F}_q))$. Then with probability at least $1 - \varepsilon$ over $b \sim \varphi_1(U_{V(\mathbb{F}_q)})$, it holds that*

$$|V_b(\mathbb{F}_q) \setminus V'_b(\mathbb{F}_q)| \leq \varepsilon \cdot |V_b(\mathbb{F}_q)|.$$

Proof. For $b \in \varphi_1(V(\mathbb{F}_q))$, let W_b be the union of the irreducible components of V_b of dimension at most $\dim V - n_1$ that are not absolutely irreducible. By Lemma 7.2, we have $|W_b(\mathbb{F}_q)| \leq d^2 q^{\dim V - n_1 - 1}$. Let $B_0 = \bigcup_{b \in \varphi_1(V(\mathbb{F}_q))} W_b(\mathbb{F}_q)$. Then

$$|B_0| \leq |\varphi_1(V(\mathbb{F}_q))| \cdot d^2 q^{\dim V - n_1 - 1} \leq d^2 q^{\dim V - 1} \leq (2d^2/q) \cdot |V(\mathbb{F}_q)|.$$

where the last inequality uses Theorem 4.6.

Let $U = \{a \in V(\mathbb{F}_q) : \dim \varphi^{-1}(\varphi(a)) = \dim V - n\}$ and $B = V(\mathbb{F}_q) \setminus U$. Then $|B| = \delta \cdot |V(\mathbb{F}_q)|$. Also let $U' = \bigcup_{b \in \varphi_1(V(\mathbb{F}_q))} V_b'(\mathbb{F}_q)$ and $B' = V(\mathbb{F}_q) \setminus U'$.

By Lemma 7.5, for every $a \in U$, the irreducible component Z of $V_{\varphi_1(a)}$ containing a satisfies that $\dim Z = \dim V - n_1$ and $\dim \overline{\varphi_2(Z)} = n_2$. So either Z is not absolutely irreducible (which implies that $a \in Z \subseteq W_{\varphi_1(a)} \subseteq B_0$), or $a \in Z \subseteq V'_{\varphi_1(a)} \subseteq U'$. It follows that $U \subseteq B_0 \cup U'$ and hence $B' \subseteq B_0 \cup B$. Therefore,

$$|B'| \leq |B_0| + |B| \leq (2d^2/q + \delta) \cdot |V(\mathbb{F}_q)| = \varepsilon^2 \cdot |V(\mathbb{F}_q)|.$$

So we have

$$\mathop{\mathbb{E}}_{b \sim \varphi_1(U_{V(\mathbb{F}_q)})} \left[\frac{|V_b(\mathbb{F}_q) \setminus V_b'(\mathbb{F}_q)|}{|V_b(\mathbb{F}_q)|} \right] = \sum_{b \in \varphi_1(V(\mathbb{F}_q))} \frac{|V_b(\mathbb{F}_q)|}{|V(\mathbb{F}_q)|} \cdot \frac{|V_b(\mathbb{F}_q) \setminus V_b'(\mathbb{F}_q)|}{|V_b(\mathbb{F}_q)|} = \frac{|B'|}{|V(\mathbb{F}_q)|} \leq \varepsilon^2.$$

By Markov's inequality, the probability that $|V_b(\mathbb{F}_q) \setminus V_b'(\mathbb{F}_q)| > \varepsilon \cdot |V_b(\mathbb{F}_q)|$ holds over $b \sim \varphi_1(U_{V(\mathbb{F}_q)})$ is at most ε. □

As a consequence of Lemma 7.6, we also prove the following lemma, which will be used in Section 9.

Lemma 7.7. *Let $V \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ be an absolutely irreducible affine variety over $\mathbb{F}_q$. Let $\varphi = (\varphi_1, \varphi_2) : V \to \mathbb{A}^n_{\mathbb{F}_q}$ be a dominant morphism defined by $f_1, \ldots, f_n$ as in Lemma 7.6. Let $d \in \mathbb{N}^+$ and $\varepsilon \in (0, 1)$ such that $\varepsilon^2 \geq 2(n + 1)d^2/q$. Assume $q \geq 20d^5$. Also assume that $f_1, \ldots, f_n \in \mathcal{L}_{h_1,\ldots,h_s,\mathbb{F}_q}$ for some $h_1, \ldots, h_s \in \mathbb{F}_q[X_1, \ldots, X_r]$ with $\deg h_1 \geq \cdots \geq \deg h_s$ such that*

$$\deg V \cdot \prod_{i=1}^{n} \deg h_i \leq d.$$

Let $D = (D_1, D_2) = \varphi(U_{V(\mathbb{F}_q)})$ where $D_i = \varphi_i(U_{V(\mathbb{F}_q)})$ for $i = 1, 2$. Then with probability at least $1 - \varepsilon$ over $b \sim D_1$, the distribution $D_2|_{D_1 = b}$ is ε-close to an (n_2, n_2, d) algebraic source over $\mathbb{F}_q$.

Proof. Note that $\deg V \leq d$ as the fact that φ is dominant implies that $\deg h_i \geq 1$ for $i \in [s]$. For $b \in \mathbb{F}_q^{n_1}$, let V_b and V_b' be as in Lemma 7.6, i.e., $V_b = \varphi_1^{-1}(b)$ and V_b' is the union of the irreducible components Z of V_b such that Z is absolutely irreducible of dimension $\dim V - n_1$ and $\dim \overline{\varphi_2(Z)} = n_2$.

Consider $a \in V(\mathbb{F}_q)$ and let $b = \varphi_1(a)$. Note that $V_b = V \cap V(f_1 - f_1(a), \ldots, f_{n_1} - f_{n_1}(a))$. Recall that the polynomial $f_1, \ldots, f_n \in \mathcal{L}_{h_1,\ldots,h_s,\mathbb{F}_q}$ are linear combinations of $h_1, \ldots, h_s$ and 1 over $\mathbb{F}_q$. By Gaussian elimination, we can find integers $1 \leq j_1 < \cdots < j_t \leq n$, where $0 \leq t \leq n_1$, and polynomials $g_1, \ldots, g_t \in \mathbb{F}_q[X_1, \ldots, X_r]$ such that $V(f_1 - f_1(a), \ldots, f_{n_1} - f_{n_1}(a)) = V(g_1, \ldots, g_t)$ and each g_i can be written as a linear combination

$$g_i = c_{i,j_i} h_{j_i} + c_{i,j_i+1} h_{j_i+1} + \cdots + c_{i,s} h_s + c_i$$

with $c_{i,j}, c_i \in \mathbb{F}_q$ and $c_{i,j_i} \neq 0$. Bézout's inequality (Lemma 4.2) then gives

$$\deg V_b \leq \deg V \cdot \prod_{i=1}^{t} \deg g_i = \deg V \cdot \prod_{i=1}^{t} \deg h_{j_i} \leq d. \tag{10}$$

Let $\{\widehat{j}_1, \dots, \widehat{j}_{s-t}\} = [s] \setminus \{j_1, \dots, j_t\}$, where $\widehat{j}_1 < \cdots < \widehat{j}_{s-t}$. As $g_1, \dots, g_t$ vanish identically on V_b, adding to each f_i a multiple of g_j for $j \in [t]$ does not change $f_i|_{V_b}$. In particular, for $i \in [n]$, we can eliminate the dependence of f_i on $h_{j_1}, \dots, h_{j_t}$ and find $\widetilde{f}_i \in \mathcal{L}_{h_{\widehat{j}_1}, \dots, h_{\widehat{j}_{s-t}}, \mathbb{F}_q}$ such that $\widetilde{f}_i|_{V_b} = f_i|_{V_b}$. Then the morphism $\varphi_2|_{V_b} : V_b \to \mathbb{A}^{n_2}_{\mathbb{F}_q}$ is defined by the polynomials $\widetilde{f}_{n_1+1}, \dots, \widetilde{f}_n$. And

$$\deg V_b' \cdot \prod_{i=1}^{n_2} \deg h_{\widehat{j}_i} \le \deg V_b \cdot \prod_{i=1}^{n_2} \deg h_{\widehat{j}_i} \overset{(10)}{\le} \deg V \cdot \prod_{i=1}^{t} \deg h_{j_i} \cdot \prod_{i=1}^{n_2} \deg h_{\widehat{j}_i} \le \deg V \cdot \prod_{i=1}^{n} \deg h_i \le d. \quad (11)$$

Let $\delta = \Pr_{a \sim U_{V(\mathbb{F}_q)}}[\dim \varphi^{-1}(\varphi(a)) \ne \dim V - n]$. By the effective fiber dimension theorem (Corollary 4.11), there exists a polynomial $P \in \overline{\mathbb{F}}_q[X_1, \dots, X_r]$ of degree at most $n \cdot \deg V \cdot \prod_{i=1}^n \deg h_i \le nd$ that does not vanish identically on $V_{\overline{\mathbb{F}}_q}$ such that for every $a \in V_{\overline{\mathbb{F}}_q}$ satisfying $P(a) \ne 0$, the fiber $\varphi^{-1}(\varphi(a))$ is equidimensional of dimension $\dim V - n$. Let B be the set of $a \in V(\mathbb{F}_q)$ such that $\dim \varphi^{-1}(\varphi(a)) \ne \dim V - n$. Then $B \subseteq V_{\overline{\mathbb{F}}_q} \cap V(P) \cap \mathbb{F}_q^r$. By Bézout's inequality (Lemma 4.2) and Lemma 4.5, we have

$$|B| \le \deg V \cdot \deg P \cdot q^{\dim V - 1} \le nd^2 q^{\dim V - 1}.$$

Let $\delta = \Pr_{a \sim U_{V(\mathbb{F}_q)}}[\dim \varphi^{-1}(\varphi(a)) \ne \dim V - n]$. Then

$$\delta = \frac{|B|}{|V(\mathbb{F}_q)|} \le \frac{nd^2 q^{\dim V - 1}}{q^{\dim V}/2} = 2nd^2/q,$$

where we use the fact $|V(\mathbb{F}_q)| \ge q^{\dim V}/2$ that follows from Theorem 4.6. So $\varepsilon \ge (2(n+1)d^2/q)^{1/2} \ge (2d^2/q + \delta)^{1/2}$.

By Lemma 7.6, with probability at least $1 - \varepsilon$ over $b \sim \varphi_1(U_{V(\mathbb{F}_q)})$, it holds that $|V_b(\mathbb{F}_q) \setminus V_b'(\mathbb{F}_q)| \le \varepsilon \cdot |V_b(\mathbb{F}_q)|$. Fix b such that this holds. Note that $D_2|_{D_1 = b} = \varphi_2(U_{V_b(\mathbb{F}_q)})$. So it suffices to verify that $\varphi_2(U_{V_b(\mathbb{F}_q)})$ is ε-close to an (n_2, n_2, d) algebraic source over $\mathbb{F}_q$.

The set $V_b'(\mathbb{F}_q)$ is nonempty as $|V_b'(\mathbb{F}_q)| \ge (1 - \varepsilon)|V_b(\mathbb{F}_q)| > 0$. By definition, every irreducible component Z of V_b' is absolutely irreducible of dimension $\dim V - n_1$ and satisfies $\dim \overline{\varphi_2(Z)} = n_2$. So the distribution $\varphi_2(U_{V_b'(\mathbb{F}_q)})$ satisfies the first two conditions of (n_2, n_2, d) algebraic sources in Definition 1.2. And the third condition also holds by (11). This shows that $\varphi_2(U_{V_b'(\mathbb{F}_q)})$ is an (n_2, n_2, d) algebraic source over $\mathbb{F}_q$.

Finally, as $|V_b(\mathbb{F}_q) \setminus V_b'(\mathbb{F}_q)| \le \varepsilon \cdot |V_b(\mathbb{F}_q)|$, the distributions $U_{V_b(\mathbb{F}_q)}$ and $U_{V_b'(\mathbb{F}_q)}$ are ε-close by Lemma 7.3. It follows that $D_2|_{D_1 = b} = \varphi_2(U_{V_b(\mathbb{F}_q)})$ is ε-close to the (n_2, n_2, d) algebraic source $\varphi_2(U_{V_b'(\mathbb{F}_q)})$. □

We now use Lemma 7.6 and the effective fiber dimension theorem to decompose irreducible (n, k, d) algebraic sources, thus completing the proof of the main result of this section.

Lemma 7.8. *Suppose $q \ge \max\{20d^5, 2(k+1)d^2/\varepsilon^2\}$, where $\varepsilon \in (0, 1)$. Then every irreducible (n, k, d) algebraic source over $\mathbb{F}_q$ is 3ε-close to a convex combination of irreducibly minimal (n, k, d) algebraic sources over $\mathbb{F}_q$.*

Proof. Let D be an irreducible (n, k, d) algebraic source over $\mathbb{F}_q$. Let $V \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ and $f : \mathbb{A}^r_{\mathbb{F}_q} \to \mathbb{A}^n_{\mathbb{F}_q}$ be as in Definition 1.2, where V is irreducible (and hence absolutely irreducible) and $D = f(U_{V(\mathbb{F}_q)})$. Let $k_V = \dim V \ge k$. By the effective Lang–Weil bound (Theorem 4.6), we have

$$|V(\mathbb{F}_q)| \ge q^{k_V}/2.$$

By Lemma 4.3, there exist distinct $j_1, \dots, j_k \in [n]$ such that the morphism $\varphi_2 : V \to \mathbb{A}^k_{\mathbb{F}_q}$ defined by $f_{j_1}, \dots, f_{j_k}$ satisfies $\dim \overline{\varphi_2(V)} = k$. View φ_2 as a morphism over $\overline{\mathbb{F}}_q$ and let $V_{\varphi_2} \subseteq \mathbb{A}^r_{\overline{\mathbb{F}}_q(Y_1,\dots,Y_k)}$ be its generic fiber. By working over the algebraically closure of $\overline{\mathbb{F}}_q(Y_1, \dots, Y_k)$ and following the proof of (10), we see that the degree of V_{φ_2} is bounded by d. As $q > 2d$, applying Lemma 4.9 with $S = \mathbb{F}_q$, we see that there exist linear polynomials $\ell_1, \dots, \ell_{k_V} \in \mathbb{F}_q[X_1, \dots, X_r]$ such that the morphisms $\pi : V \to \mathbb{A}^{k_V}_{\overline{\mathbb{F}}_q}$ defined by $\ell_1, \dots, \ell_{k_V}$ and $\tau : V_{\varphi_2} \to \mathbb{A}^{k_V - k}_{\overline{\mathbb{F}}_q(Y_1,\dots,Y_k)}$ defined by $\ell_1, \dots, \ell_{k_V - k}$ are finite.

Applying the general form of the effective fiber dimension theorem (Theorem 4.10) to φ_2, where we choose $t = k_V - k$, we see that there exists a polynomial $P \in \overline{\mathbb{F}}_q[X_1, \dots, X_r]$ of degree at most $k \cdot \deg V \cdot \prod_{i=1}^k \deg h_i \le kd$ that does not vanish identically on $V_{\overline{\mathbb{F}}_q}$ such that the following holds: Let $\varphi : V_{\overline{\mathbb{F}}_q} \to \mathbb{A}^{k_V}_{\overline{\mathbb{F}}_q}$ be the morphism defined by $\ell_1, \dots, \ell_{k_V - k}, f_{j_1}, \dots, f_{j_k}$. Then for every $a \in V_{\overline{\mathbb{F}}_q}$ satisfying $P(a) \ne 0$, it holds that $\dim \varphi^{-1}(\varphi(a)) = 0$. Note that φ is dominant by the finiteness of τ.

Let $\varphi_1 : V \to \mathbb{A}^{k_V - k}_{\mathbb{F}_q}$ be the morphism defined by $\ell_1, \dots, \ell_{k_V - k}$. View φ as a dominant morphism $V \to \mathbb{A}^{k_V}_{\mathbb{F}_q}$ over $\mathbb{F}_q$. Then $\varphi = (\varphi_1, \varphi_2)$. Let B be the set of $a \in V(\mathbb{F}_q) = V_{\overline{\mathbb{F}}_q} \cap \mathbb{F}_q^r$ satisfying $\dim \varphi^{-1}(\varphi(a)) \ne 0$. Then $B \subseteq V_{\overline{\mathbb{F}}_q} \cap V(P) \cap \mathbb{F}_q^r$. The degree of $V_{\overline{\mathbb{F}}_q} \cap V(P)$ is bounded by $\deg V \cdot \deg P \le kd^2$ by Bézout's inequality (Lemma 4.2). As P does not vanish identically on V, either $\dim(V \cap V(P)) = k_V - 1$ or $V \cap V(P) = \emptyset$. So by Lemma 4.5,

$$|B| \le |V_{\overline{\mathbb{F}}_q} \cap V(P) \cap \mathbb{F}_q^r| \le kd^2 q^{k_V - 1}.$$

Let $\delta = \Pr_{a \sim U_{V(\mathbb{F}_q)}}[\dim \varphi^{-1}(\varphi(a)) \ne 0]$. Then

$$\delta = \frac{|B|}{|V(\mathbb{F}_q)|} \le \frac{kd^2 q^{k_V - 1}}{q^{k_V}/2} = 2kd^2/q.$$

For $b \in \mathbb{F}_q^{k_V - k}$, let $V_b = \varphi_1^{-1}(b)$, and let V'_b be the union of the irreducible components Z of V_b such that Z is absolutely irreducible of dimension k and $\dim \overline{\varphi_2(Z)} = k$. As φ_1 is a linear map, we have $\deg V_b \le \deg V \le d$ by Bézout's inequality.

Note that $(2d^2/q + \delta)^{1/2} \le (2(k+1)d^2/q)^{1/2} \le \varepsilon$. By Lemma 7.6, with probability at least $1 - \varepsilon$ over $b \sim \varphi_1(U_{V(\mathbb{F}_q)})$, it holds that

$$|V_b(\mathbb{F}_q) \setminus V'_b(\mathbb{F}_q)| \le \varepsilon \cdot |V_b(\mathbb{F}_q)|. \tag{12}$$

Fix $b \in \varphi_1(V(\mathbb{F}_q))$ for which (12) holds. Then $U_{V'_b(\mathbb{F}_q)}$ is ε-close to $U_{V_b(\mathbb{F}_q)}$ by Lemma 7.3.

The set $V'_b(\mathbb{F}_q)$ is nonempty as $|V'_b(\mathbb{F}_q)| \ge (1 - \varepsilon)|V_b(\mathbb{F}_q)| > 0$. By definition, every irreducible component Z of V'_b is absolutely irreducible of dimension k and satisfies $\dim \overline{\varphi_2(Z)} = k$. This also implies that $\dim \overline{f(Z)} \ge k$ for every irreducible component Z of V'_b as the output of φ_2 is part of that of $f|_V$. So the distribution $f(U_{V'_b(\mathbb{F}_q)})$ satisfies the first two conditions of (n, k, d) algebraic sources in Definition 1.2. And the third condition also holds as $\deg V_b \le \deg V$. Finally, we know $\dim V'_b = k$. It follows that $f(U_{V'_b(\mathbb{F}_q)})$ is a minimal (n, k, d) algebraic source over $\mathbb{F}_q$. Therefore, $f(U_{V_b(\mathbb{F}_q)})$ is ε-close to a minimal (n, k, d) algebraic source over $\mathbb{F}_q$.

Let $D' = U_{V(\mathbb{F}_q)}$ so that $D = f(D')$. For each $b \in \varphi_1(V(\mathbb{F}_q))$, the distribution $D|_{\varphi_1(D') = b}$ is exactly $f(U_{V_b(\mathbb{F}_q)})$. We have already shown that with probability at least $1 - \varepsilon$ over $b \sim \varphi_1(U_{V(\mathbb{F}_q)})$, the distribution $D|_{\varphi_1(D')=b} = f(U_{V_b(\mathbb{F}_q)})$ is ε-close to the minimal (n,k,d) algebraic source $f(U_{V'_b(\mathbb{F}_q)})$. It follows that D is 2ε-close to a convex combination of minimal (n,k,d) algebraic source over $\mathbb{F}_q$.

Finally, by Lemma 7.4, every minimal (n,k,d) algebraic source over $\mathbb{F}_q$ is ε-close to a convex combination of irreducibly minimal (n,k,d) algebraic sources over $\mathbb{F}_q$. It follows that D is 3ε-close to a convex combination of irreducibly minimal (n,k,d) algebraic sources over $\mathbb{F}_q$. □

Combining Lemma 7.4 and Lemma 7.8 yields the following corollary.

Corollary 7.9 (Decomposition into irreducibly minimal algebraic sources)**.** *Suppose* $q \geq \max\{20d^5, 2(k+1)d^2/\varepsilon^2\}$*, where* $\varepsilon \in (0,1)$*. Then every* (n,k,d) *algebraic source over* $\mathbb{F}_q$ *is* 4ε*-close to a convex combination of irreducibly minimal* (n,k,d) *algebraic sources over* $\mathbb{F}_q$*.*

7.2 Estimating the Min-Entropy of (n,k,d) Algebraic Sources

We first prove the following lower bound on the min-entropy of an (n,k,d) algebraic source D (or more precisely, a distribution D' close to D). The proof uses the decomposition into irreducible (n,k,d) algebraic sources (Lemma 7.4).

Lemma 7.10. *Suppose* $q \geq \max\{20d^5, 2kd^2/\varepsilon\}$*, where* $\varepsilon \in (0,1/2]$*. Then every* (n,k,d) *algebraic source over* $\mathbb{F}_q$ *is* 2ε*-close to a* k'*-source over the set* $\mathbb{F}_q^n$*, where* $k' = k\log q - \log d - 2$*.*

Proof. Assume $k > 0$ as otherwise the lemma holds trivially. Let D be an (n,k,d) algebraic source over $\mathbb{F}_q$. By Lemma 7.4, we know D is ε-close to an irreducible (n,k,d) algebraic source D' over $\mathbb{F}_q$. Suppose $D' = f(U_{V(\mathbb{F}_q)})$ where $V \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ and $f : \mathbb{A}^r_{\mathbb{F}_q} \to \mathbb{A}^n_{\mathbb{F}_q}$ are as in Definition 1.2, and V is absolutely irreducible. So f is defined by polynomials $f_1, \dots, f_n \in \mathcal{L}_{h_1,\dots,h_s,\mathbb{F}_q}$, where $h_1, \dots, h_s \in \mathbb{F}_q[X_1, \dots, X_r]$, $\deg h_1 \geq \cdots \geq \deg h_s$, and $\deg V \cdot \prod_{i=1}^k \deg h_i \leq d$. By the effective Lang–Weil bound (Theorem 4.6), we have $|V(\mathbb{F}_q)| \geq q^{\dim V}/2$.

By Lemma 4.3, there exist distinct $i_1, \dots, i_k \in [n]$ such that the polynomial map $\psi : \mathbb{A}^r_{\mathbb{F}_q} \to \mathbb{A}^k_{\mathbb{F}_q}$ defined by $f_{i_1}, \dots, f_{i_k}$ satisfies $\dim \overline{\psi(V)} = k$. Applying the effective fiber dimension theorem (Corollary 4.11) to ψ (viewed as a morphism over $\overline{\mathbb{F}}_q$), we see that there exists a polynomial $P \in \overline{\mathbb{F}}_q[X_1, \dots, X_r]$ of degree at most $k \cdot \deg V \cdot \prod_{i=1}^k \deg h_i \leq kd$ that does not vanish identically on $V_{\overline{\mathbb{F}}_q}$ such that for every $a \in V_{\overline{\mathbb{F}}_q}$ satisfying $P(a) \neq 0$, it holds that $\dim \psi|^{-1}_{V_{\overline{\mathbb{F}}_q}}(\psi(a)) = \dim V - k$.

Let $B = \{a \in V(\mathbb{F}_q) : P(a) = 0\} = V_{\overline{\mathbb{F}}_q} \cap V(P) \cap \mathbb{F}_q^r$. It follows from Lemma 4.5 and Bézout's inequality (Lemma 4.2) that

$$|B| \leq \deg V \cdot \deg P \cdot q^{\dim V - 1} \leq kd^2 q^{\dim V - 1}.$$

Let U' be the uniform distribution over $V(\mathbb{F}_q) \setminus B$. By Lemma 7.3, the distributions $U_{V(\mathbb{F}_q)}$ and U' are ε'-close, where

$$\varepsilon' = \frac{|B|}{|V(\mathbb{F}_q)|} \leq \frac{kd^2 q^{\dim V - 1}}{q^{\dim V}/2} = 2kd^2/q \leq \varepsilon.$$

So $D' = f(U_{V(\mathbb{F}_q)})$ and $f(U')$ are ε-close. It follows that D and $f(U')$ are 2ε-close.

It remains to prove that $f(U')$ has min-entropy at least $k\log q - \log d - 2$. As $\psi(U')$ can be obtained from $f(U')$ by projecting to a subset of coordinates, it suffices to show that $\psi(U')$ has min-entropy at least $k\log q - \log d - 2$. Consider arbitrary $a \in V(\mathbb{F}_q) \setminus B$ and let $b = \psi(a)$. We have $P(a) \neq 0$ and hence $\dim \psi|^{-1}_V(b) = \dim V - k$. The fiber $\psi|^{-1}_V(b)$ is the subvariety of V defined by the k polynomials $f_{i_1} - f_{i_1}(a), \dots, f_{i_k} - f_{i_k}(a) \in \mathcal{L}_{h_1,\dots,h_s,\mathbb{F}_q}$. By Gaussian elimination, we can construct polynomials $g_1, \ldots, g_t$ from $f_{i_1} - f_{i_1}(a), \ldots, f_{i_k} - f_{i_k}(a)$ such that $t \leq k$, $\psi|_V^{-1}(b) = V \cap V(g_1, \ldots, g_t)$, and $g_i \in \mathcal{L}_{h_i, \ldots, h_s, \mathbb{F}_q}$ for $i \in [t]$. In particular, we have $\deg g_i \leq \deg h_i$ for $i \in [t]$. It follows from Bézout's inequality (Lemma 4.2) that

$$\deg \psi|_V^{-1}(b) \leq \deg V \cdot \prod_{i=1}^{t} \deg g_i \leq \deg V \cdot \prod_{i=1}^{k} \deg h_i \leq d.$$

Lemma 4.5 then gives $\left|\left(\psi|_V^{-1}(b)\right)(\mathbb{F}_q)\right| \leq dq^{\dim V - k}$. Therefore,

$$\Pr[\psi(U') = b] = \frac{\left|\left(\psi|_V^{-1}(b)\right)(\mathbb{F}_q)\right|}{|V(\mathbb{F}_q) \setminus B|} \leq \frac{\left|\left(\psi|_V^{-1}(b)\right)(\mathbb{F}_q)\right|}{|V(\mathbb{F}_q)|/2} \leq \frac{dq^{\dim V - k}}{q^{\dim V}/4} = 4d/q^k.$$

Every element in the support of $\psi(U')$ has the form $b = \psi(a)$ for some $a \in V(\mathbb{F}_q) \setminus B$. So $\psi(U')$ has min-entropy at least $-\log(4d/q^k) = k \log q - \log d - 2$, as desired. □

The next proposition complements Lemma 7.10 and gives an upper bound on the min-entropy.

Proposition 7.11. *Suppose $q \geq 20d^5$. Let D be an (n, k, d) algebraic source over $\mathbb{F}_q$ such that k is maximal with respect to this condition, i.e., D is not an $(n, k+1, d)$ algebraic source over $\mathbb{F}_q$. Then the statistical distance between D and any $(k \log q + 2 \log d + 2)$-source is at least $\frac{1}{4d}$. Moreover, if D is an irreducible (n, k, d) algebraic source over $\mathbb{F}_q$, then the statistical distance between D and any $(k \log q + \log d + 1)$-source is at least $\frac{1}{2}$.*

Proof. Suppose $D = f(U_{V(\mathbb{F}_q)})$ where $V \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ and $f : \mathbb{A}^r_{\mathbb{F}_q} \to \mathbb{A}^n_{\mathbb{F}_q}$ are as in Definition 1.2. As D is not an $(n, k+1, d)$ algebraic source over $\mathbb{F}_q$, we know V has an irreducible component V_0 of dimension $\dim V$ that is absolutely irreducible such that the dimension of $\overline{f(V_0)}$ is exactly k. We have $|V_0(\mathbb{F}_q)| \geq q^{\dim V}/2$ by the effective Lang–Weil bound (Theorem 4.6) and the assumption that $q \geq 20d^5$. Also note that $|V(\mathbb{F}_q)| \leq dq^{\dim V}$ by Lemma 4.5.

Let $W = \overline{f(V_0)}$. Then $\deg W \leq d$ by Lemma 4.12 and the third condition in Definition 1.2. So $|W(\mathbb{F}_q)| \leq dq^k$ by Lemma 4.5.

Let D' be a k'-source over the set $\mathbb{F}_q^n$, where $k' = k \log q + 2 \log d + 2$. Then

$$\Pr[D' \in W(\mathbb{F}_q)] \leq |W(\mathbb{F}_q)| \cdot 2^{-k'} \leq dq^k 2^{-k'} = \frac{1}{4d}.$$

On the other hand, as $f(V_0(\mathbb{F}_q)) \subseteq W(\mathbb{F}_q)$ and $D = f(U_{V(\mathbb{F}_q)})$, we have

$$\Pr[D \in W(\mathbb{F}_q)] \geq \frac{|V_0(\mathbb{F}_q)|}{|V(\mathbb{F}_q)|} \geq \frac{q^{\dim V}/2}{dq^{\dim V}} = \frac{1}{2d}.$$

So the statistical distance between D and D' is at least $\frac{1}{2d} - \frac{1}{4d} = \frac{1}{4d}$.

Now assume D is an irreducible (n, k, d) algebraic source over $\mathbb{F}_q$. So V may be chosen to be irreducible. Then $V_0 = V$ and hence $\Pr[D \in W(\mathbb{F}_q)] = 1$. Let $k'' = k \log q + \log d + 1$ and let D'' be a k''-source over the set $\mathbb{F}_q^n$. Then $\Pr[D'' \in W(\mathbb{F}_q)] \leq |W(\mathbb{F}_q)| 2^{-k''} \leq dq^k 2^{-k''} = \frac{1}{2}$. So the statistical distance between D and D'' is at least $1 - \frac{1}{2} = \frac{1}{2}$. □

8 Extracting a Short Seed

In this section, we consider the problem of constructing explicit deterministic extractors for (n, k, d) algebraic sources over a finite field $\mathbb{F}_q$ in the special case where $k = 1$.

The main results of this section are explicit constructions of deterministic extractors that extract almost $\log q$ bits from $(1, 1, d)$ algebraic sources and, more generally, $(n, 1, d)$ algebraic sources over $\mathbb{F}_q$. They are used as building blocks in the construction of the full-fledged deterministic extractors that extract most min-entropy from (n, k, d) algebraic sources.

Formally, we prove the following theorems.

Theorem 8.1 (Extractor for $(1, 1, d)$ algebraic sources)**.** *Let $d \in \mathbb{N}^+$ and $\varepsilon \in (0, 1/2]$. Suppose $q \geq c_0 d^5/\varepsilon^2$, where $c_0 > 0$ is a large enough absolute constant. Then there exists an explicit ε-extractor $\mathsf{Ext} : \mathbb{F}_q \to \{0, 1\}^m$ for $(1, 1, d)$ algebraic sources over $\mathbb{F}_q$ such that $m \geq \log q - 2 \log \log p - O(\log(d/\varepsilon))$.*

Theorem 8.2 (Extractor for $(n, 1, d)$ algebraic sources)**.** *Let $d \in \mathbb{N}^+$ and $\varepsilon \in (0, 1/2]$. Suppose $q \geq (nd/\varepsilon)^{c_0}$, where $c_0 > 0$ is a large enough absolute constant. Then there exists an explicit ε-extractor $\mathsf{Ext} : \mathbb{F}_q \to \{0, 1\}^m$ for $(n, 1, d)$ algebraic sources over $\mathbb{F}_q$ such that $m \geq \log q - 2 \log \log p - O(\log(nd/\varepsilon))$.*

Theorem 8.2 is derived from Theorem 8.1. As in [DGW09, Dvi12], the proof of Theorem 8.1 uses Bombieri's estimate for exponential sums (Theorem 4.7). However, the argument in [DGW09, Dvi12] works only when the characterisitic p is large. Moreover, it only yields an extractor that extracts $c \log q$ bits for some constant $c \leq 1/2$. We introduce new ideas that allow us to extract almost $\log q$ bits regardless of the characteristic p.

8.1 Exponential Sums over Finite Fields of Arbitrary Characteristic

The purpose of this subsection is to prove the following estimate for exponential sums over curves, even over finite fields of small characteristics. Recall that Bombieri's estimate (Theorem 4.7) is valid as long as the polynomial f does not have the form $g^p - g$ on the curve. One way to deal with this difficulty is to require p to be large. However, we would like to get meaningful results for arbitrary p, and we do this by paying the cost of excluding a small subgroup of characters from the estimate.

Lemma 8.3. *Let $C \subseteq \mathbb{A}^n_{\mathbb{F}_q}$ be an irreducible affine curve of degree d_1 over a finite field $\mathbb{F}_q$ of characteristic p, and let $f \in \mathbb{F}_q[X_1, \ldots, X_n]$ be a polynomial of degree d_2 that is not constant on C. Then the set of characters $\chi \in \widehat{\mathbb{F}_q}$ for which*

$$\left| \sum_{x \in C(\mathbb{F}_q)} \chi(f(x)) \right| \leq (d_1^2 + 2d_1 d_2 - 3d_1) q^{1/2} + d_1^2 \tag{13}$$

fails to hold is contained in a subgroup of $\widehat{\mathbb{F}_q}$ of size at most $d_1 d_2$.

Before proving Lemma 8.3, we require some preliminary results. Recall that for a formal Laurent series $f \in \mathbb{F}((T))$, we denote by $\mathrm{ord}(f)$ the least degree of the terms that appear in f, i.e., $f = c_0 T^{\mathrm{ord}(f)} + c_1 T^{\mathrm{ord}(f)+1} + \cdots$ where $c_0 \neq 0$. And $\mathrm{ord}(f) = +\infty$ if $f = 0$. For $i \in \mathbb{Z}$, denote by $\mathrm{coeff}_i(f)$ the coefficient of T^i in f.

Lemma 8.4. *Let $C_0 \subseteq \mathbb{A}^n$ be an irreducible affine curve of degree d over an algebraically closed field $\mathbb{F}$. Then there exists an $\mathbb{F}$-linear field embedding $\tau : \mathbb{F}(C_0) \hookrightarrow \mathbb{F}((T))$ such that for any polynomial $f \in \mathbb{F}[X_1, \dots, X_n]$ of degree d that is not constant on C_0, the map τ sends f to $\tilde{f} \in \mathbb{F}((T))$ such that*

$$-\deg(C_0) \cdot d \le \operatorname{ord}(\tilde{f}) < 0.$$

We defer the proof of Lemma 8.4 to Appendix C.

Lemma 8.5. *Let $\mathbb{F}$ be a field of characteristic p. Suppose $f, g \in \mathbb{F}((T))$ such that $f = g^p - g$ and* $\operatorname{ord}(f) < 0$*. Let $t < 0$ and $e \ge 0$ be integers such that t is coprime to p and $tp^{e+1} < \operatorname{ord}(f) \le tp^e$. Then*

$$\sum_{i=0}^{e} (\operatorname{coeff}_{tp^i}(f))^{p^{e-i}} = 0.$$

Proof. The lemma can be proved by expressing the coefficients of f in terms of those of g. We give an alternative proof. Note that by taking a field extension, we may assume $\mathbb{F}$ is algebraically closed and hence a perfect field. Then the map $a \mapsto a^{1/p}$ is an automorphism of $\mathbb{F}$. We will instead prove

$$\sum_{i=0}^{+\infty} (\operatorname{coeff}_{tp^i}(f))^{1/p^i} = 0. \tag{14}$$

Note that the LHS of (14) is actually a finite sum and equals $\sum_{i=0}^{e} (\operatorname{coeff}_{tp^i}(f))^{1/p^i}$. The lemma then follows by raising both sides to the p^e-th power.

Note that to prove (14), we may assume all terms in g have degree tp^i for $i \in \mathbb{N}$, as ignoring the other terms does not affect the coefficients appearing in (14). Moreover, as $\operatorname{coeff}_{tp^i}(\cdot)$ and the map $x \mapsto x^{1/p}$ are both linear in characteristic p, if (14) holds for $f_1, f_2 \in \mathbb{F}((T))$, then it also holds for $f_1 + f_2$. So we may assume g contains only one term cT^{tp^i} where $c \in \mathbb{F}$ and $i \in \mathbb{N}$, and hence $f = c^p T^{tp^{i+1}} - cT^{tp^i}$. The LHS of (14) is then $(c^p)^{1/p^{i+1}} - c^{1/p^i} = 0$. □

Proof of Lemma 8.3. Fix an irreducible component C_0 of the affine curve $C_{\overline{\mathbb{F}}_q}$ over $\overline{\mathbb{F}}_q$. It is a standard fact that the natural inclusion $\mathbb{F}_q[C] \hookrightarrow \overline{\mathbb{F}}_q[C_{\overline{\mathbb{F}}_q}]$ induces an inclusion $\mathbb{F}_q[C] \hookrightarrow \overline{\mathbb{F}}_q[C_0]$.[5] In particular, as f is not constant on C, it is not constant on C_0 either.

Denote by σ be the character $x \mapsto e^{2\pi i x/p}$ of $\mathbb{F}_p$. For $\alpha \in \mathbb{F}_q$, denote by χ_α the character of $\mathbb{F}_q$ sending x to $(\sigma \circ \operatorname{Tr})(\alpha x)$. The map $\alpha \mapsto \chi_\alpha$ is a one-to-one correspondence between $\mathbb{F}_q$ and $\widehat{\mathbb{F}_q}$. Consider a character χ_α for which (13) does not hold. Then by Bombieri's estimate (Theorem 4.7), there exists $g \in \overline{\mathbb{F}}_q[X_1, \dots, X_n]$ such that $\alpha f - (g^p - g)$ vanishes identically on C, and hence also on C_0.

View f and g as elements of $\overline{\mathbb{F}}_q(C_0)$. By Lemma 8.4, there exists an $\overline{\mathbb{F}}_q$-linear field embedding $\tau : \overline{\mathbb{F}}_q(C_0) \hookrightarrow \overline{\mathbb{F}}_q((T))$ such that $\tilde{f} := \tau(f)$ satisfies

$$-\deg(C_0) \cdot d_2 \le \operatorname{ord}(\tilde{f}) < 0.$$

[5] This uses the fact that $\overline{\mathbb{F}}_q[C_{\overline{\mathbb{F}}_q}]$ is an *integral extension* of $\mathbb{F}_q[C]$. The kernel of $\mathbb{F}_q[C] \to \overline{\mathbb{F}}_q[C_0]$ is $I \cap \mathbb{F}_q[C]$, where I is the minimal prime ideal of $\overline{\mathbb{F}}_q[C_{\overline{\mathbb{F}}_q}]$ defining the irreducible component C_0. As $C_{\overline{\mathbb{F}}_q}$ is equidimensional of dimension one, I is not a maximal ideal of $\overline{\mathbb{F}}_q[C_{\overline{\mathbb{F}}_q}]$. Then $I \cap \mathbb{F}_q[C]$ is a prime ideal of $\mathbb{F}_q[C]$ that is not maximal either by [AM69, Corollary 5.8]. As C is an irreducible curve, this implies $I \cap \mathbb{F}_q[C] = 0$. See [AM69, Chapter 5] for more details about integral extensions.

Let $\tilde{g} = \tau(g)$. Then $\alpha\tilde{f} = \tilde{g}^p - \tilde{g}$ as τ is an $\overline{\mathbb{F}}_q$-linear field embedding and $\alpha f = g^p - g$ in $\overline{\mathbb{F}}_q[C_0]$. Write $\mathrm{ord}(\tilde{f}) = tp^e$ where $t < 0$ is coprime to p. By Lemma 8.5,

$$0 = \sum_{i=0}^{e} (\mathrm{coeff}_{tp^i}(\alpha\tilde{f}))^{p^{e-i}} = \sum_{i=0}^{e} \alpha^{p^{e-i}} (\mathrm{coeff}_{tp^i}(\tilde{f}))^{p^{e-i}}. \tag{15}$$

The RHS of (15) is a nonzero polynomial in α independent of g, and its degree is bounded by p^e. Also note that this polynomial is a linearized polynomial, i.e., the degree of every monomial is a power of p. So its roots in $\mathbb{F}_q$ form a subgroup H whose size is at most p^e.

Then the set $S := \{\chi_\alpha : \alpha \in H\}$ contains all the characters for which (13) fails to hold. As χ_α is defined by $x \mapsto (\sigma \circ \mathrm{Tr})(\alpha x)$ and $\mathrm{Tr}(\cdot)$ is linear, the set S is a subgroup of $\widehat{\mathbb{F}_q}$, whose size is at most $|H| \le p^e \le |\mathrm{ord}(\tilde{f})| \le \deg(C_0) \cdot d_2 \le d_1 d_2$. □

Remark. We need the curve C to be irreducible in Lemma 8.3 so that the "bad" characters are contained in a single subgroup of $\widehat{\mathbb{F}_q}$ of size at most $d_1 d_2$. For a reducible curve C, a similar proof shows that these characters are contained in a subset $S \subseteq \widehat{\mathbb{F}_q}$ of size at most $d_1 d_2$ such that S is the union of a collection of subgroups of $\widehat{\mathbb{F}_q}$, one for each irreducible component of C.

8.2 Proofs of Theorem 8.1 and Theorem 8.2

Theorem 8.1 can be proved easily using the techniques we have developed so far. First, we show that a distribution of the form $f(U_{C(\mathbb{F}_q)})$ is a strongly (ε, d)-biased source, and even an ε-biased source if the characteristic p is large, where C is a low-degree absolutely irreducible affine curve over $\mathbb{F}_q$ and f is a low-degree polynomial.

Lemma 8.6. *Let $C \subseteq \mathbb{A}^n_{\mathbb{F}_q}$ be an absolutely irreducible affine curve of degree d_1 over a finite field $\mathbb{F}_q$ of characteristic p. Let $f \in \mathbb{F}_q[X_1, \dots, X_n]$ be a polynomial of degree d_2 that is not constant on C. Let $d \ge d_1 d_2$ and suppose $q \ge 20d^5$. Then $f(U_{C(\mathbb{F}_q)})$ is strongly (ε, d)-biased over the set $\mathbb{F}_q$, where $\varepsilon = 8d^2/q^{1/2}$. Furthermore, if $p > d_1 d_2$, then $f(U_{C(\mathbb{F}_q)})$ is ε-biased over the set $\mathbb{F}_q$.*

Proof. By Lemma 8.3, there exists a subgroup $S \subseteq \widehat{\mathbb{F}_q}$ of size most $d_1 d_2 \le d$ such that for all $\chi \in \widehat{\mathbb{F}_q} \setminus S$, it holds that

$$\left| \sum_{x \in C(\mathbb{F}_q)} \chi(f(x)) \right| \le (d_1^2 + 2d_1 d_2 - 3d_1) q^{1/2} + d_1^2 \le 4d^2 q^{1/2}.$$

By the effective Lang–Weil bound (Theorem 4.6) and the assumption $q \ge 20d^5$, we have $|C(\mathbb{F}_q)| \ge q/2$. It follows that for $\chi \in \widehat{\mathbb{F}_q} \setminus S$,

$$\left|\mathbb{E}[\chi(f(U_{C(\mathbb{F}_q)}))]\right| = \frac{\left|\sum_{x \in C(\mathbb{F}_q)} \chi(f(x))\right|}{|C(\mathbb{F}_q)|} \le \frac{4d^2 q^{1/2}}{q/2} = 8d^2/q^{1/2} = \varepsilon.$$

By definition, this means $f(U_{C(\mathbb{F}_q)})$ is strongly (ε, d)-biased.

Now assume $p > d_1 d_2$. As the size of $S \subseteq \widehat{\mathbb{F}_q}$ is a power of p and $|S| \le d_1 d_2$, we must have $|S| = 1$, i.e., S contains only the trivial character. So $f(U_{C(\mathbb{F}_q)})$ is ε-biased. □

We are now ready to prove Theorem 8.1. This follows by first reducing to the case of irreducibly minimal $(1,1,d)$ algebraic sources and then applying Lemma 8.6 together with the constructions in Section 3.

Proof of Theorem 8.1. We will construct an $(\varepsilon/2)$-extractor Ext for irreducibly minimal $(1,1,d)$ algebraic sources over $\mathbb{F}_q$ with the claimed output length. By Corollary 7.9 and the assumption that $q \geq c_0 d^5/\varepsilon^2$, every $(1,1,d)$ algebraic source over $\mathbb{F}_q$ is $(\varepsilon/2)$-close to a convex combination of irreducibly minimal $(1,1,d)$ algebraic sources over $\mathbb{F}_q$. It follows that Ext is also an ε-extractor for $(1,1,d)$ algebraic sources over $\mathbb{F}_q$.

Let $\varepsilon_0 := 8d^2/q^{1/2}$ and let $c > 0$ be a large enough absolute constant. We consider two different cases depending on how large the characteristic p is.

Case 1: $p \leq (d/\varepsilon)^c$. In this case, let $\mathsf{Ext} : \mathbb{F}_q \to \mathbb{F}_p^t = \{0,1\}^m$ be the $(\varepsilon/2)$-extractor for strongly (ε_0, d)-biased sources given by Theorem 3.21, where we set the parameters

$$\begin{aligned} n' &= \min\{\lfloor 2\log_p(1/\varepsilon_0) - 2\log_p(16d/(\varepsilon/2)^2)\rfloor, \log_p q\}, \\ t &= \lfloor n' - 3 - 2\log_p(2d/(\varepsilon/2))\rfloor, \text{ and} \\ m &= t \log p. \end{aligned}$$

As $\varepsilon_0 = 8d^2/q^{1/2}$ and $p \leq (d/\varepsilon)^c$, we have

$$m \geq \min\{2\log(1/\varepsilon_0), \log q\} - O(\log(d/\varepsilon)) - O(\log p) = \log q - O(\log(d/\varepsilon)).$$

Consider an irreducibly minimal $(1,1,d)$ algebraic source D over $\mathbb{F}_q$. By definition, there exist $r \in \mathbb{N}^+$, an absolutely irreducible affine curve $C \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ over $\mathbb{F}_q$ of degree d_1, and a polynomial $f \in \mathbb{F}_q[X_1, \ldots, X_r]$ of degree d_2 that is not constant on C such that $d_1 d_2 \leq d$ and $D = f(U_{C(\mathbb{F}_q)})$. By Lemma 8.6, the distribution D is a strongly (ε_0, d)-biased distribution over the set $\mathbb{F}_q$, so that $\mathsf{Ext}(D)$ is $(\varepsilon/2)$-close to the uniform distribution over $\mathbb{F}_q$. It follows that Ext is an $(\varepsilon/2)$-extractor for irreducibly minimal $(1,1,d)$ algebraic sources, as desired.

Case 2: $p > (d/\varepsilon)^c \geq d_1 d_2$. In this case, identify $\mathbb{F}_q$ with the abelian group $\mathbb{Z}_N^t$, where $N = p$ and $t = \log_p q$. Let Ext be the map $\mathbb{Z}_N^t \to \mathbb{Z}_N^{t-1} \times \mathbb{Z}_M$ in Lemma 3.9, where the parameter M will be determined shortly. By Lemma 3.9, Ext is an ε'-extractor for ε_0-biased distribution, where

$$\varepsilon' := \varepsilon_0 \cdot (N^{t-1}M)^{1/2} \cdot C \log N + M/N = 8d^2 \cdot (M/p)^{1/2} \cdot C \log p + M/p.$$

and $C > 0$ is an absolute constant. By Lemma 8.6 and the fact that $p > d_1 d_2$, every irreducibly minimal $(1,1,d)$ algebraic source over $\mathbb{F}_q$ is ε_0-biased. So Ext is also an ε'-extractor for irreducibly minimal $(1,1,d)$ algebraic sources.

We want to choose $M \in \mathbb{N}^+$ such that $\varepsilon' \leq \varepsilon/2$. As $p > (d/\varepsilon)^c$ and $c > 0$ is a large enough constant, such an integer M exists and we can choose M such that $\log M \geq \log p - 2\log\log p - O(\log(d/\varepsilon))$. So Ext is an $(\varepsilon/2)$-extractor for irreducibly minimal $(1,1,d)$ algebraic sources with the output length

$$m = (t-1)\log N + \log M = \log q - \log p + \log M \geq \log q - 2\log\log p - O(\log(d/\varepsilon)).$$

In both cases above, we ignore the technicality that the size of the range of Ext may not be a power of two, i.e., m may not be an integer. But one can always turn the size into a power of two at the cost of losing $O(\log(1/\varepsilon))$ bits in the output by composing Ext with a suitable map. The details are left to the reader. □

Remark. In the case where $p \le (d/\varepsilon)^c$, the above proof shows that we could avoid losing $2\log\log p$ bits in the output. However, this does not make an essential difference as $2\log\log p$ is dominated by the term $O(\log(d/\varepsilon))$ in this case.

We now prove Theorem 8.2 by composing the extractors in Theorem 8.1 with the deterministic rank extractors for varieties constructed in Section 6.

Proof of Theorem 8.2. Choose sufficiently large $d' = \Theta(nd^2 \log n)$. Let $\mathsf{Ext}' : \mathbb{F}_q \to \{0,1\}^m$ be an explicit $(\varepsilon/2)$-extractor for $(1,1,d')$ algebraic sources over $\mathbb{F}_q$ as constructed in Theorem 8.1, where

$$m \ge \log q - 2\log\log p - O(\log(d'/\varepsilon)) = \log q - 2\log\log p - O(\log(nd/\varepsilon)).$$

Let $\varphi : \mathbb{A}^n_{\overline{\mathbb{F}}_q} \to \mathbb{A}^1_{\overline{\mathbb{F}}_q}$ be an explicit $(n,1,d)$ deterministic rank extractor for varieties defined by a polynomial $F \in \mathbb{F}_q[X_1, \dots, X_n]$ as constructed in Corollary 6.5, where $\deg F = O(nd \log n)$. View φ as a morphism $\mathbb{A}^n_{\mathbb{F}_q} \to \mathbb{A}^1_{\mathbb{F}_q}$ over $\mathbb{F}_q$.

Let $\mathsf{Ext} := \mathsf{Ext}' \circ \varphi|_{\mathbb{F}_q^n} : \mathbb{F}_q^n \to \{0,1\}^m$. We claim that Ext is an ε-extractor for $(n,1,d)$ algebraic sources over $\mathbb{F}_q$. To see this, consider an irreducibly minimal $(n,1,d)$ algebraic source $D = f(U_{C(\mathbb{F}_q)})$ arising from an absolutely irreducible affine curve $C \subseteq \mathbb{A}^n_{\mathbb{F}_q}$ and a polynomial map $f : \mathbb{A}^r_{\mathbb{F}_q} \to \mathbb{A}^n_{\mathbb{F}_q}$ defined by polynomials $f_1, \dots, f_n$. Let $d_1 = \deg C$ and $d_2 = \max\{\deg f_1, \dots, \deg f_n\}$. We have $d_1 d_2 \le d$ by Definition 1.2. Then $\varphi \circ f : \mathbb{A}^r_{\mathbb{F}_q} \to \mathbb{A}^1_{\mathbb{F}_q}$ is defined by the polynomial $F(f_1, \dots, f_n)$ of degree $O(d_2 \cdot nd \log n)$. Note that $\deg C \cdot \deg F(f_1, \dots, f_n) = O(d_1 d_2 \cdot nd \log n) = O(nd^2 \log n) \le d'$. So $\varphi(D) = (\varphi \circ f)(U_{C(\mathbb{F}_q)})$ satisfies the third condition of $(1,1,d')$ algebraic sources over $\mathbb{F}_q$ in Definition 1.2 with respect to C and $\varphi \circ f$.

We have $\dim \overline{f(C)} = 1$ by Definition 1.2 and $\deg \overline{f(C)} \le d_1 d_2 \le d$ by Lemma 4.12. As φ is an $(n,1,d)$ deterministic rank extractor for varieties, we have $\dim \overline{(\varphi \circ f)(C)} = 1$. It follows that $\varphi(D) = (\varphi \circ f)(U_{C(\mathbb{F}_q)})$ is a $(1,1,d')$ algebraic source over $\mathbb{F}_q$. As Ext' is an explicit $(\varepsilon/2)$-extractor for $(1,1,d')$ algebraic sources over $\mathbb{F}_q$, $\mathsf{Ext}(D) = \mathsf{Ext}'(\varphi(D))$ is $(\varepsilon/2)$-close to U_m.

The above proof shows that Ext is an $(\varepsilon/2)$-extractor for irreducibly minimal $(n,1,d)$ algebraic sources over $\mathbb{F}_q$. By Corollary 7.9, every $(n,1,d)$ algebraic sources over $\mathbb{F}_q$ is $(\varepsilon/2)$-close to a convex combination of irreducibly minimal $(n,1,d)$ algebraic sources over $\mathbb{F}_q$. So Ext is an ε-extractor for $(n,1,d)$ algebraic sources over $\mathbb{F}_q$. □

9 Deterministic Extractors for (n,k,d) Algebraic Sources

In this section, we provide our main construction of deterministic extractors for (n,k,d) algebraic sources. Recall that in Section 8 we considered the case of $(n,1,d)$ algebraic sources.

We start with the case of (n,n,d) algebraic sources, and we follow our general proof technique as laid out in Section 1.3: the first step of the construction is applying our extractor from Section 8 to obtain a short output, which is then, in the second step, used as a seed for a seeded extractor for sources with high min-entropy (note that even though we have more structure in our source, since we are anyway applying a seeded extractor we might as well use an off-the-shelf construction which works for any source with high min-entropy). Proving that this indeed works requires analyzing the conditional distribution of an (n,n,d) algebraic source under fixing of a subset of the coordinates, which is done in Lemma 9.1. This construction is presented and analyzed in Section 9.1.

In order to remove the assumption that $k = n$ and handle general (n,k,d) algebraic sources, we apply a rank extractor which, roughly speaking, condenses a k-dimensional source in an ambient

n-dimensional space to a k-dimensional source in an ambient k-dimensional space, and this enables us to use the extractor from Section 9.1. As discussed at the end of Section 9.1, this can be done using the deterministic rank extractor of Section 6, but it would have an undesirable effect on the field size. Thus, we opt to use a *linear seeded rank extractor* (as defined in Section 5), where the seed of the rank extractor is chosen pseudorandomly using our extractor for $(n, 1, d)$ algebraic sources from Section 8.

To summarize, in our composition theorem (Theorem 9.6), we start by applying the extractor for $(n, 1, d)$ algebraic sources from Section 8 in order to select a seed for the seeded linear rank extractor from Section 5, we apply the resulting linear map to the source, and then we use the extractor for full-rank sources from Section 9.1 to obtain the final output. The details of this construction appear in Section 9.2.

9.1 Deterministic Extractors for Full-Rank Algebraic Sources

The following lemma states that irreducible (n, n, d) algebraic sources have a nice recursive structure. The statement can be extended to general (n, k, d) algebraic sources in some way, but this special case is simpler and suffices for us.

Lemma 9.1. *Suppose $q \geq \max\{20d^5, 2(n+1)d^2/\varepsilon^2\}$, where $\varepsilon \in (0, 1)$. Let $D = (D_1, D_2)$ be an irreducible (n, n, d) algebraic source over $\mathbb{F}_q$, where D_1 and D_2 are distributions over $\mathbb{F}_q^{n_1}$ and $\mathbb{F}_q^{n_2}$ respectively and $n_1 + n_2 = n$. Then the following holds:*

1. *D_1 is an irreducible (n_1, n_1, d) algebraic source over $\mathbb{F}_q$.*
2. *With probability at least $1-\varepsilon$ over $b \sim D_1$, the distribution $D_2|_{D_1=b}$ is ε-close to an (n_2, n_2, d) algebraic source over $\mathbb{F}_q$.*

Proof. Suppose $D = f(U_{V(\mathbb{F}_q)})$ where $V \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ and $f : \mathbb{A}^r_{\mathbb{F}_q} \to \mathbb{A}^n_{\mathbb{F}_q}$ are as in Definition 1.2. So V is absolutely irreducible and $\dim \overline{f(V)} = n$. And f is defined by polynomials $f_1, \ldots, f_n \in \mathcal{L}_{h_1,\ldots,h_s,\mathbb{F}_q}$, where $h_1, \ldots, h_s \in \mathbb{F}_q[X_1, \ldots, X_r]$, $\deg h_1 \geq \cdots \geq \deg h_s$, and $\deg V \cdot \prod_{i=1}^n \deg h_i \leq d$.

Let $\varphi_1 : V \to \mathbb{A}^{n_1}_{\mathbb{F}_q}$ be the morphism defined by $f_1, \ldots, f_{n_1}$, and similarly, let $\varphi_2 : V \to \mathbb{A}^{n_2}_{\mathbb{F}_q}$ be the morphism defined by $f_{n_1+1}, \ldots, f_n$. Then $f|_V = (\varphi_1, \varphi_2)$ and $D_i = \varphi_i(U_{V(\mathbb{F}_q)})$ for $i = 1, 2$.

We know V is absolutely irreducible. And the dimension of $\overline{\varphi_1(V)}$ must be n_1 since otherwise $\dim \overline{f(V)}$ cannot reach n. By definition, D_1 is an irreducible (n_1, n_1, d) algebraic source over $\mathbb{F}_q$. This proves the first claim.

As $\varepsilon^2 \geq 2(n+1)d^2/q$ and $q \geq 20d^5$, by Lemma 7.7, with probability at least $1-\varepsilon$ over $b \sim D_1$, the distribution $D_2|_{D_1=b}$ is ε-close to an (n_2, n_2, d) algebraic source over $\mathbb{F}_q$, proving the second claim. □

We also need the following explicit construction of seeded extractors given by Goldreich and Wigderson [GW97], which is based on expander graphs.

Theorem 9.2 ([GW97])**.** *For $n \in \mathbb{N}$, $0 \leq \Delta \leq n$ and $\varepsilon > 0$, there exists an explicit seeded ε-extractor $\mathsf{Ext} : \{0,1\}^n \times \{0,1\}^\ell \to \{0,1\}^n$ for $(n-\Delta)$-sources with $\ell = O(\Delta + \log(1/\varepsilon))$.*

We now provide our construction for full-rank algebraic sources. Our construction follows the general paradigm mentioned in Section 1.3: we first apply our extractor from Theorem 8.1 to obtain a short output, which is then used as a seed to the extractor from Theorem 9.2. Proving that this "randomness recycling" technique works in this setting requires Lemma 9.1.

Theorem 9.3 (Extractor for (n,n,d) algebraic sources)**.** *Let $n, d \in \mathbb{N}^+$ and $\varepsilon \in (0,1/2]$. Suppose $q \geq (nd/\varepsilon)^{c_0}$, where $c_0 > 0$ is a large enough absolute constant. Then there exists an explicit ε-extractor $\mathsf{Ext} : \mathbb{F}_q \to \{0,1\}^m$ for (n,n,d) algebraic sources over $\mathbb{F}_q$ such that $m \geq n \log q - 2 \log\log p - O(\log(d/\varepsilon))$.*

Proof. If $n = 1$, then the statement holds by Theorem 8.1. So assume $n > 1$. By Corollary 7.9, every (n,n,d) algebraic source over $\mathbb{F}_q$ is $(\varepsilon/2)$-close to an irreducible (n,n,d) algebraic source over $\mathbb{F}_q$. So it suffices to construct an explicit $(\varepsilon/2)$-extractor Ext for irreducible (n,n,d) algebraic sources over $\mathbb{F}_q$ with the claimed output length.

Let $\varepsilon' = \varepsilon/10$. We construct Ext as follows.

1. Let $m_1 = \lceil (n-1)\log q \rceil$ and $\Delta = \log d + 3$. Let $\mathsf{Ext}_1 : \mathbb{F}_q^{n-1} \times \{0,1\}^\ell \to \{0,1\}^{m_1}$ be an explicit seeded ε'-extractor for k-sources, where $k = m_1 - \Delta \leq (n-1)\log q - \log d - 2$ and $\ell = O(\Delta + \log(1/\varepsilon))$. This can be done by using Theorem 9.2 to construct an ε'-extractor $\mathsf{Ext}_1' : \{0,1\}^{m_1} \times \{0,1\}^\ell \to \{0,1\}^{m_1}$ for k-sources and then composing it with an injection $\mathbb{F}_q^{n-1} \hookrightarrow \{0,1\}^{m_1}$.

2. Let $\mathsf{Ext}_2 : \mathbb{F}_q \to \{0,1\}^{m_2}$ be an explicit ε'-extractor for $(1,1,d)$ algebraic sources over $\mathbb{F}_q$ such that $m_2 \geq \log q - 2\log\log p - O(\log(d/\varepsilon))$. Moreover, we assume the constant $c_0 > 0$ is large enough so that $m_2 \geq \ell = O(\log(d/\varepsilon))$. Such an extractor can be constructed by Theorem 8.1.

3. The map Ext takes $(x_1, x_2) \in \mathbb{F}_q^{n-1} \times \mathbb{F}_q$ and feeds x_2 to Ext_2. Let $y = (y_1, y_2)$ be the output of Ext_2, where $y_1 \in \{0,1\}^\ell$. (This is possible as $m_2 \geq \ell$.) Then Ext outputs $(\mathsf{Ext}_1(x_1, y_1), y_2)$.

Ext outputs $m := m_1 + m_2 - \ell$ bits. Plugging in the values of m_1, m_2, ℓ. We see that the output length of Ext is as claimed.

Let D be an irreducibe (n,n,d) algebraic sources over $\mathbb{F}_q$. Write $D = (D_1, D_2)$ where D_1 is distributed over $\mathbb{F}_q^{n-1}$ and D_2 is distributed over $\mathbb{F}_q$.

By Lemma 9.1, with probability at least $1-\varepsilon'$ over $x_1 \sim D_1$, the distribution $D_2|_{D_1 = x_1}$ is ε'-close to a $(1,1,d)$ algebraic source over $\mathbb{F}_q$. As Ext_2 is an ε'-extractor for $(1,1,d)$ algebraic sources over $\mathbb{F}_q$, we see that with probability at least $1-\varepsilon'$ over $x_1 \sim D_1$, it holds that $\mathsf{Ext}_2(D_2)|_{D_1=x_1} =_{2\varepsilon'} U_{m_2}$. By Lemma 2.2,

$$(D_1, \mathsf{Ext}_2(D_2)) =_{3\varepsilon'} D_1 \times U_{m_2}.$$

By Lemma 9.1, D_1 is an $(n-1, n-1, d)$ algebraic source over $\mathbb{F}_q$. By Lemma 7.10 and the fact that $k \leq (n-1)\log q - \log d - 2$, the distribution D_1 is ε'-close to a k-source D_1'. So

$$(D_1, \mathsf{Ext}_2(D_2)) =_{4\varepsilon'} D_1' \times U_{m_2}. \tag{16}$$

By the definition of Ext, (16) implies

$$\mathsf{Ext}(D) =_{4\varepsilon'} \mathsf{Ext}_1(D_1' \times U_\ell) \times U_{m_2 - \ell}. \tag{17}$$

As Ext_1 is a seeded ε'-extractor for k-sources, we see

$$\mathsf{Ext}_1(D_1' \times U_\ell) =_{\varepsilon'} U_{m_1}. \tag{18}$$

It follows from (17) and (18) that $\mathsf{Ext}(D) =_{5\varepsilon'} U_{m_1} \times U_{m_2-\ell} = U_m$. As $5\varepsilon' = \varepsilon/2$, we see that Ext is an $(\varepsilon/2)$-extractor for irreducible (n,n,d) algebraic sources over $\mathbb{F}_q$, and hence an ε-extractor for (n,n,d) algebraic sources over $\mathbb{F}_q$. □

One can remove the full-rank assumption and construct an extractor for (n, k, d) algebraic sources over $\mathbb{F}_q$ by composing the extractor in Theorem 9.3 with the deterministic rank extractor for varieties in Section 6. This argument was used by Dvir, Gabizon and Wigderson [DGW09], except that they considered polynomial sources only and used a different construction of deterministic rank extractors. The downside of this argument, however, is that such a deterministic rank extractor is necessarily nonlinear. In particular, our rank extractor uses polynomials of degree at least $\mathrm{poly}(n)$, and so does the one in [DGW09]. Composing with such a rank extractor increases the degree of each polynomial in the polynomial map by at least a $\mathrm{poly}(n)$ factor. The resulting field size q would then depend at least polynomially on n^k, or n^n if $k = \Theta(n)$, assuming that we want to extract about $k \log q$ bits.

In the next subsection, we show how to remove the full-rank assumption more efficiently using a linear seeded rank extractor for varieties.

9.2 Removing the Full-Rank Assumption

We now remove the full-rank assumption in Theorem 9.3 without significantly increasing the required field size. This is done by extending an argument in [GRS06, GR08].

Lemma 9.4 ([GRS06, Lemma 2.6]). *Let $D = (D_1, D_2)$ be a joint distribution over a finite product set $A \times B$. Suppose D is ε-close to $U_A \times D_2$. Then for all $y \in \mathrm{supp}(D_1)$, the conditional distribution $D_2|_{D_1=y}$ is ε'-close to D_2, where $\varepsilon' = 2\varepsilon \cdot |A|$.*

We also need the following lemma.

Lemma 9.5. *Suppose $q \geq \max\{20d^5, 2(k+1)d^2/\varepsilon^2\}$, where $\varepsilon \in (0, 1)$. Let $D = f(U_{V(\mathbb{F}_q)})$ be an irreducible (n, k, d) algebraic source over $\mathbb{F}_q$ arising from an affine variety $V \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ and a polynomial map $f : \mathbb{A}^r_{\mathbb{F}_q} \to \mathbb{A}^n_{\mathbb{F}_q}$ as in Definition 1.2. Let $\pi : \mathbb{A}^n_{\mathbb{F}_q} \to \mathbb{A}^{k-1}_{\mathbb{F}_q}$ be a linear map over $\mathbb{F}_q$ such that $\dim \overline{(\pi \circ f)(V)} = k-1$. Then with probability at least $1-\varepsilon$ over $b \sim \pi(D)$, the distribution $D|_{\pi(D)=b}$ is ε-close to an $(n, 1, d)$ algebraic source over $\mathbb{F}_q$.*

Proof. The proof is similar to that of Lemma 7.7, although we are not able to derive the statement directly from Lemma 7.7 for technical reasons.

By definition, V is absolutely irreducible, and f is defined by polynomials $f_1, \dots, f_n \in \mathcal{L}_{h_1,\dots,h_s,\mathbb{F}_q}$, where $h_1, \dots, h_s \in \mathbb{F}_q[X_1, \dots, X_r]$, $\deg h_1 \geq \cdots \geq \deg h_s$, and $\deg V \cdot \prod_{i=1}^k \deg h_i \leq d$.

Append a coordinate Y_u of $\mathbb{A}^n_{\mathbb{F}_q}$ for some $u \in [n]$ to the output of π to obtain a linear map $\pi' : \mathbb{A}^n_{\mathbb{F}_q} \to \mathbb{A}^k_{\mathbb{F}_q}$ such that $\dim \overline{(\pi' \circ f)(V)} = k$. This is possible by Lemma 4.3.

Let $\varphi = (\pi' \circ f)|_V : V \to \mathbb{A}^k_{\mathbb{F}_q}$. Then φ is dominant. We can write $\varphi = (\varphi_1, \varphi_2)$ where $\varphi_1 : V \to \mathbb{A}^{k-1}_{\mathbb{F}_q}$ equals $(\pi \circ f)|_V$ and $\varphi_2 : V \to \mathbb{A}^1_{\mathbb{F}_q}$ is defined by the polynomial f_u. As π is a linear map over $\mathbb{F}_q$ and the set $\mathcal{L}_{h_1,\dots,h_s,\mathbb{F}_q}$ is closed under taking linear combinations over $\mathbb{F}_q$, we know φ is defined by polynomials in $\mathcal{L}_{h_1,\dots,h_s,\mathbb{F}_q}$.

Let $\delta = \Pr_{a \sim U_{V(\mathbb{F}_q)}}[\dim \varphi^{-1}(\varphi(a)) \neq \dim V - k]$. We first bound δ. By the effective fiber dimension theorem (Corollary 4.11), there exists a polynomial $P \in \overline{\mathbb{F}}_q[X_1, \dots, X_r]$ of degree at most $k \cdot \deg V \cdot \prod_{i=1}^k \deg h_i \leq kd$ that does not vanish identically on $V_{\overline{\mathbb{F}}_q}$ such that for every $a \in V_{\overline{\mathbb{F}}_q}$ satisfying $P(a) \neq 0$, the fiber $\varphi^{-1}(\varphi(a))$ is equidimensional of dimension $\dim V - k$. Let B be the set of $a \in V(\mathbb{F}_q)$ such that $\dim \varphi^{-1}(\varphi(a)) \neq \dim V - k$. Then $B \subseteq V_{\overline{\mathbb{F}}_q} \cap V(P) \cap \mathbb{F}_q^r$. By Bézout's inequality (Lemma 4.2) and Lemma 4.5, we have

$$|B| \le \deg V \cdot \deg P \cdot q^{\dim V - 1} \le kd^2 q^{\dim V - 1}.$$

Therefore,

$$\delta = \frac{|B|}{|V(\mathbb{F}_q)|} \le \frac{kd^2 q^{\dim V - 1}}{q^{\dim V}/2} = 2kd^2/q,$$

where we use the fact $|V(\mathbb{F}_q)| \ge q^{\dim V}/2$ that follows from Theorem 4.6. So $\varepsilon \ge (2(k+1)d^2/q)^{1/2} \ge (2d^2/q + \delta)^{1/2}$.

For $b \in \mathbb{F}_q^{k-1}$, let $V_b = \varphi_1^{-1}(b)$ and let V_b' be the union of the irreducible components Z of V_b such that Z is absolutely irreducible of dimension $\dim V - (k-1)$ and $\dim \overline{\varphi_2(Z)} = 1$. By Lemma 7.6, with probability at least $1 - \varepsilon$ over $b \sim \varphi_1(U_{V(\mathbb{F}_q)}) = \pi(D)$, it holds that

$$|V_b(\mathbb{F}_q) \setminus V_b'(\mathbb{F}_q)| \le \varepsilon \cdot |V_b(\mathbb{F}_q)|. \tag{19}$$

Fix b such that (19) holds. Note that $D|_{\pi(D)=b} = f(U_{V_b(\mathbb{F}_q)})$. By (19) and Lemma 7.3, the distributions $f(U_{V_b(\mathbb{F}_q)})$ and $f(U_{V_b'(\mathbb{F}_q)})$ are ε-close. So it suffices to verify that $f(U_{V_b'(\mathbb{F}_q)})$ is an $(n, 1, d)$ algebraic source over $\mathbb{F}_q$.

The set $V_b'(\mathbb{F}_q)$ is nonempty as $|V_b'(\mathbb{F}_q)| \ge (1-\varepsilon)|V_b(\mathbb{F}_q)| > 0$. By definition, every irreducible component Z of V_b' is absolutely irreducible of dimension $\dim V - (k-1)$ and satisfies $\dim \overline{\varphi_2(Z)} = 1$. This also implies that $\dim \overline{f(Z)} \ge 1$ for every irreducible component Z of V_b' as φ_2 is defined by f_u and hence its output is part of that of $f|_V$. So the distribution $f(U_{V_b'(\mathbb{F}_q)})$ satisfies the first two conditions of $(n, 1, d)$ algebraic sources in Definition 1.2.

We now verify the third condition in Definition 1.2. As in the proof of Lemma 7.7, by Gaussian elimination, we can find integers $1 \le j_1 < \cdots < j_t \le n$, where $0 \le t \le k-1$, and polynomials $g_1, \ldots, g_t \in \mathbb{F}_q[X_1, \ldots, X_r]$ such that $V_b = V \cap V(g_1, \ldots, g_t)$, and each g_i can be written as a linear combination

$$g_i = c_{i,j_i} h_{j_i} + c_{i,j_i+1} h_{j_i+1} + \cdots + c_{i,s} h_s + c_i$$

with $c_{i,j}, c_i \in \mathbb{F}_q$ and $c_{i,j_i} \ne 0$. Bézout's inequality (Lemma 4.2) then gives

$$\deg V_b \le \deg V \cdot \prod_{i=1}^{t} \deg g_i = \deg V \cdot \prod_{i=1}^{t} \deg h_{j_i}. \tag{20}$$

Let $\{\widehat{j}_1, \ldots, \widehat{j}_{s-t}\} = [s] \setminus \{j_1, \ldots, j_t\}$, where $\widehat{j}_1 < \cdots < \widehat{j}_{s-t}$. As $g_1, \ldots, g_t$ vanish identically on V_b, adding to each f_i a multiple of g_j for $j \in [t]$ does not change $f_i|_{V_b}$. In particular, for $i \in [n]$, we can eliminate the dependence of f_i on $h_{j_1}, \ldots, h_{j_t}$ and find $\widetilde{f}_i \in \mathcal{L}_{h_{\widehat{j}_1}, \ldots, h_{\widehat{j}_{s-t}}, \mathbb{F}_q}$ such that $\widetilde{f}_i|_{V_b} = f_i|_{V_b}$. Then the morphism $f|_{V_b} : V_b \to \mathbb{A}^n_{\mathbb{F}_q}$ is defined by the polynomials $\widetilde{f}_1, \ldots, \widetilde{f}_n$. And

$$\deg V_b' \cdot \deg h_{\widehat{j}_1} \le \deg V_b \cdot \deg h_{\widehat{j}_1} \overset{(20)}{\le} \deg V \cdot \left(\prod_{i=1}^{t} \deg h_{j_i} \right) \cdot \deg h_{\widehat{j}_1} \le \deg V \cdot \prod_{i=1}^{k} \deg h_i \le d.$$

So the third condition in Definition 1.2 is also satisfied (with respect to the polynomials $\widetilde{f}_i$). This shows that $f(U_{V_b'(\mathbb{F}_q)})$ is an $(n, 1, d)$ algebraic source over $\mathbb{F}_q$, as desired. □

The following theorem shows how to compose all the ingredients in our construction: an extractor Ext_1 for $(n,1,d)$ algebraic sources, an extractor Ext_2 for full-rank algebraic sources, and a linear seeded rank extractor φ, in order to obtain extractors for (n,k,d) algebraic sources. The construction uses Ext_1 in order to select the seed for φ, applies φ on the input, and then applies Ext_2 on the resulting “condensed” source.

Theorem 9.6 (Composition of extractors)**.** *Let $n \geq k > 1$ be integers. Let $\varepsilon, \varepsilon' \in (0,1)$. Suppose we are given the following objects:*

- *an ε-extractor $\mathsf{Ext}_1 : \mathbb{F}_q^n \to \{0,1\}^{m_1}$ for $(n,1,d)$ algebraic sources over $\mathbb{F}_q$,*
- *an ε-extractor $\mathsf{Ext}_2 : \mathbb{F}_q^{k-1} \to \{0,1\}^{m_2}$ for $(k-1,k-1,d)$ algebraic sources over $\mathbb{F}_q$, and*
- *an $(n,k-1,k,\varepsilon')$ linear seeded rank extractor $(\varphi_y)_{y\in\{0,1\}^\ell}$ for varieties over $\overline{\mathbb{F}}_q$ (see Definition 5.1) such that $\ell \leq m_1$ and each φ_y is defined by linear polynomials over $\mathbb{F}_q$.*

Write $\mathsf{Ext}_1 = (\mathsf{Ext}_1', \mathsf{Ext}_1'')$, where Ext_1' and Ext_1'' output the first ℓ bits and the last $m_1 - \ell$ bits of Ext_1 respectively. Assume $q \geq \max\{20d^5, 2(k+1)d^2/\varepsilon^2\}$. Then the map $\mathsf{Ext} : \mathbb{F}_q^n \to \{0,1\}^{m_1} \times \{0,1\}^{m_2} = \{0,1\}^{m_1+m_2}$ defined by

$$\mathsf{Ext}(x) := (\mathsf{Ext}_1(x), \mathsf{Ext}_2(\varphi_{\mathsf{Ext}_1'(x)}(x)))$$

is a $(6\varepsilon \cdot 2^\ell + 4\varepsilon + \varepsilon')$-extractor for (n,k,d) algebraic sources over $\mathbb{F}_q$.

Towards the proof of Theorem 9.6, we start with the following lemma.

Lemma 9.7. *Use the notations and assumptions in Theorem 9.6. Let D be an irreducible (n,k,d) algebraic source over $\mathbb{F}_q$ arising from an affine variety $V \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ and a polynomial map $f : \mathbb{A}^r_{\mathbb{F}_q} \to \mathbb{A}^n_{\mathbb{F}_q}$ as in Definition 1.2. Let $y \in \operatorname{supp}(\mathsf{Ext}_1'(D))$. Assume $\dim \overline{\varphi_y(f(V))} = k-1$. Then*

$$(\mathsf{Ext}_1''(D), \mathsf{Ext}_2(\varphi_y(D)))|_{\mathsf{Ext}_1'(D)=y} =_{3\varepsilon\cdot 2^\ell+\varepsilon} U_{m_1+m_2-\ell}.$$

Proof. By Lemma 9.5, with probability at least $1-\varepsilon$ over $b \sim \varphi_y(D)$, the distribution $D|_{\varphi_y(D)=b}$ is ε-close to an $(n,1,d)$ algebraic source D_b' over $\mathbb{F}_q$. Fix $b \in \operatorname{supp}(\varphi_y(D))$ such that this happens, i.e.,

$$D|_{\varphi_y(D)=b} =_\varepsilon D_b'. \tag{21}$$

As Ext_1 is an ε-extractor for $(n,1,d)$ algebraic sources over $\mathbb{F}_q$, we have

$$\mathsf{Ext}_1(D_b') =_\varepsilon U_{m_1}. \tag{22}$$

Combining (21) and (22), we conclude that with probability at least $1-\varepsilon$ over $b \sim \varphi_y(D)$,

$$\mathsf{Ext}_1(D)|_{\varphi_y(D)=b} =_{2\varepsilon} U_{m_1}.$$

By Lemma 2.2, this implies

$$(\mathsf{Ext}_1'(D), \mathsf{Ext}_1''(D), \varphi_y(D)) = (\mathsf{Ext}_1(D), \varphi_y(D)) =_{3\varepsilon} U_{m_1} \times \varphi_y(D).$$

Therefore,

$$(\mathsf{Ext}_1'(D), \mathsf{Ext}_1''(D), \mathsf{Ext}_2(\varphi_y(D))) =_{3\varepsilon} U_{m_1} \times \mathsf{Ext}_2(\varphi_y(D)) = U_\ell \times U_{m_1-\ell} \times \mathsf{Ext}_2(\varphi_y(D)).$$

By Lemma 9.4,

$$(\mathsf{Ext}''_1(D), \mathsf{Ext}_2(\varphi_y(D))|_{\mathsf{Ext}'_1(D)=y} =_{6\varepsilon\cdot 2^\ell} U_{m_1-\ell} \times \mathsf{Ext}_2(\varphi_y(D)). \tag{23}$$

By assumption, V is irreducible and $\dim \overline{\varphi_y(f(V))} \geq k-1$. As φ_y is a linear map over $\mathbb{F}_q$ and the polynomials that define f are in $\mathcal{L}_{h_1,\dots,h_s,\mathbb{F}_q}$, which is closed under taking linear combinations over $\mathbb{F}_q$, we see that $\varphi_y(D) = (\varphi_y \circ f)(U_{V(\mathbb{F}_q)})$ is a $(k-1, k-1, d)$ algebraic source over $\mathbb{F}_q$. As Ext_2 is an ε-extractor for $(k-1, k-1, d)$ algebraic source over $\mathbb{F}_q$, we have

$$\mathsf{Ext}_2(\varphi_y(D)) =_\varepsilon U_{m_2}. \tag{24}$$

Combining (23) and (24) yields

$$(\mathsf{Ext}''_1(D), \mathsf{Ext}_2(\varphi_y(D))|_{\mathsf{Ext}'_1(D)=y} =_{6\varepsilon\cdot 2^\ell+\varepsilon} U_{m_1+m_2-\ell}$$

as desired. □

Proof of Theorem 9.6. Let D be an irreducible (n, k, d) algebraic source over $\mathbb{F}_q$ arising from an affine variety $V \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ and a polynomial map $f : \mathbb{A}^r_{\mathbb{F}_q} \to \mathbb{A}^n_{\mathbb{F}_q}$ as in Definition 1.2. Then $\dim \overline{f(V)} \geq k$. Let T be the set of $y \in \{0,1\}^\ell$ such that $\dim \overline{\varphi_y(f(V))} \geq k-1$. As $(\varphi_y)_{y\in\{0,1\}^\ell}$ is an $(n, k-1, k, \varepsilon')$ linear seeded rank extractor, we have

$$\Pr_{y\sim U_\ell}[y \in T] \geq 1 - \varepsilon'. \tag{25}$$

As D is an (n, k, d) algebraic source and hence an $(n, 1, d)$ algebraic source over $\mathbb{F}_q$, and Ext_1 is an ε-extractor for $(n, 1, d)$ algebraic source over $\mathbb{F}_q$, we have $\mathsf{Ext}_1(D) =_\varepsilon U_{m_1}$. So $\mathsf{Ext}'_1(D) =_\varepsilon U_\ell$. Combining this with (25) yields

$$\Pr_{y\sim \mathsf{Ext}'_1(D)}[y \in T] \geq 1 - \varepsilon - \varepsilon'. \tag{26}$$

By Lemma 9.7, we have

$$(\mathsf{Ext}''_1(D), \mathsf{Ext}_2(\varphi_y(D)))|_{\mathsf{Ext}'_1(D)=y} =_{6\varepsilon\cdot 2^\ell+\varepsilon} U_{m_1+m_2-\ell} \quad \text{for all } y \in \operatorname{supp}(\mathsf{Ext}'_1(D)) \cap T. \tag{27}$$

By Lemma 2.2, (26) and (27) together yield

$$\mathsf{Ext}(D) = (\mathsf{Ext}'_1(D), \mathsf{Ext}''_1(D), \mathsf{Ext}_2(\varphi_{\mathsf{Ext}'_1(D)}(D))) =_{6\varepsilon\cdot 2^\ell+2\varepsilon+\varepsilon'} \mathsf{Ext}'_1(D) \times U_{m_1+m_2-\ell}.$$

Using the fact $\mathsf{Ext}'_1(D) =_\varepsilon U_\ell$ again, we obtain

$$\mathsf{Ext}(D) =_{6\varepsilon\cdot 2^\ell+3\varepsilon+\varepsilon'} U_{m_1+m_2}.$$

The above proof shows that Ext is a $(6\varepsilon \cdot 2^\ell + 3\varepsilon + \varepsilon')$-extractor for irreducible (n, k, d) algebraic sources over $\mathbb{F}_q$. By Lemma 7.4, every (n, k, d) algebraic source over $\mathbb{F}_q$ is ε-close to a convex combination of irreducible (n, k, d) algebraic sources over $\mathbb{F}_q$. So Ext is a $(6\varepsilon \cdot 2^\ell + 4\varepsilon + \varepsilon')$-extractor for (n, k, d) algebraic sources over $\mathbb{F}_q$. □

Putting it together. We now instantiate the objects in Theorem 9.6 and prove Theorem 1.

Proof of Theorem 1. If $k = 0$, the theorem holds trivially, and we may even use an extractor with an empty output. If $k = 1$, the theorem follows from Theorem 8.2. So assume $k > 1$.

Let $\ell = \lceil \log(2n^2/\varepsilon) \rceil$. Construct an explicit linear $(n, k-1, k-1, \varepsilon_0)$ seeded rank extractor $(\varphi_y)_{y \in \{0,1\}^\ell}$ for varieties with $\varepsilon_0 = (k-1)(n-k+1)/2^\ell \leq \varepsilon/2$ using the construction in Lemma 5.2 (see also Corollary 5.4). This is possible as $|\mathbb{F}_q^\times| = q - 1 \geq \max\{n, 2^\ell\}$. By definition, $(\varphi_y)_{y \in \{0,1\}^\ell}$ is also a linear $(n, k-1, k, \varepsilon_0)$ seeded rank extractor for varieties.

Let $\varepsilon_1 = \frac{\varepsilon/2}{6 \cdot 2^\ell + 4}$. Construct an explicit ε_1-extractor $\mathsf{Ext}_1 : \mathbb{F}_q^n \to \{0,1\}^{m_1}$ for $(n, 1, d)$ algebraic sources over $\mathbb{F}_q$ using Theorem 8.2 such that

$$m_1 \geq \log q - 2 \log\log p - O(\log(nd/\varepsilon_1)) = \log q - 2 \log\log p - O(\log(nd/\varepsilon)).$$

We may assume $m_1 \geq \ell$ as $q \geq (nd/\varepsilon)^c$ where $c > 0$ is a sufficiently large constant. Write $\mathsf{Ext}_1 = (\mathsf{Ext}_1', \mathsf{Ext}_1'')$, where Ext_1' and Ext_1'' output the first ℓ bits and the last $m_1 - \ell$ bits of Ext_1 respectively.

Finally, construct an explicit ε_1-extractor $\mathsf{Ext}_2 : \mathbb{F}_q^{k-1} \to \{0,1\}^{m_2}$ for $(k-1, k-1, d)$ algebraic sources over $\mathbb{F}_q$ using Theorem 9.3 such that

$$m_2 \geq (k-1) \log q - 2 \log\log p - O(\log(d/\varepsilon_1)) = (k-1) \log q - 2 \log\log p - O(\log(nd/\varepsilon)).$$

The choice of ε_0 and ε_1 guarantees that $6\varepsilon_1 \cdot 2^\ell + 4\varepsilon_1 + \varepsilon_0 \leq \varepsilon$. By Theorem 9.6, the map $\mathsf{Ext} : \mathbb{F}_q^n \to \{0,1\}^{m_1+m_2}$ defined by $\mathsf{Ext}(x) := (\mathsf{Ext}_1(x), \mathsf{Ext}_2(\varphi_{\mathsf{Ext}_1'(x)}(x)))$ is an ε-extractor for (n, k, d) algebraic sources over $\mathbb{F}_q$, whose output length is

$$m_1 + m_2 \geq k \log q - 4 \log\log p - O(\log(nd/\varepsilon))$$

as desired. □

Remark. If $d = 1$, an (n, k, d) algebraic source over $\mathbb{F}_q$ is simply an affine source over $\mathbb{F}_q$. In this case, our output length in Theorem 1 is $k \log q - O(\log\log p + \log(n/\varepsilon))$, which is slightly better than the output length $(k-1) \log q$ in [GR08], This is due to a more careful analysis that we use. Namely, we use the fact that a linear seeded rank extractor is *strong* in the sense that most seeds are good. This allows us to include the seed in the output, which yields the improved output length. We remark that the analysis of [GR08] can be easily modified to achieve such an improvement too. Finally, when $d = 1$, an $(k-1, k-1, d)$ algebraic source is already the uniform distribution, as observed in [GR08]. So one can simply use the identity map $\mathbb{F}_q^{k-1} \to \mathbb{F}_q^{k-1}$ as Ext_2 and get a slightly better parameter $m_2 = (k-1) \log q$ instead of $m_2 = (k-1) \log q - 2 \log\log p - O(\log(n/\varepsilon))$.

10 Affine Extractors with Exponentially Small Error for Quasipolynomially Large Fields

In this section, we construct affine extractors with exponentially small error, over prime fields of size $q = n^{O(\log\log(n))}$ and any characteristic. Our construction is in fact identical to the extractor of Bourgain, Dvir and Leeman [BDL16], but our analysis is slightly improved. Specifically, Bourgain, Dvir and Leeman constructed an affine extractor over prime fields $\mathbb{F}_q$ where $q = n^{O(\log\log n)}$ is a so-called "typical" prime. Our construction works over any prime finite field of the same size.

Recall that "log" denotes logarithms in base 2. We use "ln" in this section to denote natural logarithms.

Definition 10.1 (Divisor counting function)**.** *For positive integer $n \geq 1$, let $\omega(n)$ count the number of distinct prime factors of n, not counting multiplicity.*

Bourgain, Dvir and Leeman use average-case bounds on $\omega(q-1)$ for a "typical" prime q. For our purposes we need the following worst-case upper bound on $\omega(n)$.

Lemma 10.2 ([Rob83], Theorem 13)**.** *Let $n \geq 26$ be an integer. Then*

$$\omega(n) \leq \frac{\ln n}{\ln \ln n - 1.1714}$$

The following proposition replaces the use of [BDL16] by finding a set of degrees $d_1, \dots, d_n$ with useful properties for the construction.

Proposition 10.3. *Let q be a prime number. Fix $\varepsilon > 0$. Then, if $q \geq n^{\frac{2}{\varepsilon}\log\log(n)}$, there exists an efficient deterministic algorithm that, in time polynomial in n, finds n integers $d_1 < d_2 < \cdots < d_n \in \mathbb{N}$ such that $LCM(d_1, \dots, d_n) \leq q^{\varepsilon}$ and each d_i is coprime to $q-1$.*

Notice that these properties of $d_1, \dots, d_n$ are precisely those needed for the affine extractor for [BDL16].

Proof. Suppose that $q = n^{C\log\log(n)}$ for a constant $C > 0$ to be specified later. Let $r = \lceil \log n \rceil$. Let $p_1, \dots, p_r$ be the least r primes that are coprime to $q-1$. Then $p_1, \dots, p_r$ all belong to the first $\omega(q-1) + r$ primes.

By Lemma 10.2, we know that for large enough n,

$$\omega(q-1) \leq \frac{\ln q}{\ln \ln q - 1.1714} \leq 1.01 \frac{\log q}{\log\log q}.$$

(Recall that we use log for logarithm in base 2). Therefore, $p_1, \dots, p_r$ are among the first m primes, where

$$\begin{aligned} m &\leq \lceil \log n \rceil + 1.01 \cdot \frac{\log(q)}{\log\log(q)} \\ &\leq 1 + \log(n) + \frac{1.01C\log(n) \cdot \log\log(n)}{\log(C) + \log\log(n) + \log\log\log(n)} \\ &\leq 1 + \log(n) + 1.01C\log(n) \leq 2C\log(n) \end{aligned}$$

The last two inequalities are for large enough n and C.

Next, we would like to bound the magnitude of these m primes. By the Chebyshev bound on the prime counting function (see, e.g., Theorem 5.4 in [Sho09]), $[t]$ contains at least $\frac{t\ln(2)}{2\ln(t)}$ primes. Therefore, taking $t = C'\log(n) \cdot \log\log(n)$, where $C' = 10C$, we get that for large enough n, $[t]$ contains at least m primes.

Let $D = p_1 p_2 \cdots p_r$. Notice that D has $2^{\lceil \log n \rceil} \geq n$ distinct divisors, each of which is coprime to $q-1$. These distinct divisors are the $d_1, \dots, d_n$ in the theorem statement. In particular, we have $\mathrm{LCM}(d_1, \dots, d_n) \leq D$.

We can upper-bound D as p_r^r. We want to choose q large enough that $D \leq q^{\varepsilon}$. Since $p_r \leq t$ and $r \leq \log(n) + 1$, we obtain

$$D \leq (C'\log(n) \cdot \log\log(n))^{\log(n)+1}$$

$$\begin{aligned}&\leq n^{\log(C'\log(n)\cdot\log\log(n))+1}\\&= n^{\log\log(n)+\log\log\log(n)+\log(C')+1}\end{aligned}$$

Moreover, $q^{\varepsilon} = n^{C\varepsilon\cdot\log\log(n)}$. For large enough n, it follows that choosing $C \geq \frac{2}{\varepsilon}$ is enough to imply that $D \leq q^{\varepsilon}$.

Finally, observe that indeed $d_1, \ldots, d_n$ can be found in time which is polynomial in n, by checking all integers up to $t = O(\log n \log\log n)$ for primality and co-primality with $q-1$ in order to compute $p_1, \ldots, p_r$, and then multiplying all non-empty subsets of $p_1, \ldots, p_r$ to output $d_1, \ldots, d_n$ □

The rest of the proof continues in a very similar manner to the proof of Theorem 3.1 of [BDL16]. For completeness, we provide the main details of the construction and its proof. The following lemma from [BDL16] gives a convenient form for representing affine subspaces.

Lemma 10.4 ([BDL16], Lemma 3.4)**.** *Let $V \subseteq \mathbb{F}^n$ be an affine subspace of dimension k. Then, there is an affine map $\ell : \mathbb{F}^k \to \mathbb{F}^n$ whose image is V such that there exist k indices $1 \leq j_1 < j_2 < \cdots < j_k \leq n$, such that*

1. *For all $i < j_1$, $\ell_i(t)$ is a constant in $\mathbb{F}_q$.*
2. *For all $i \in [k]$, $\ell_{j_i}(t) = t_i$.*
3. *For every $i > 1$ and $j < j_i$, $\ell_j(t)$ is an affine function which depends only on $t_1, \ldots, t_{i-1}$.*

We also need the following exponential sum estimate due to Deligne (see [Del74, MK93, BDL16]). For $b \in \mathbb{F}_q^n$, define $\chi_b : \mathbb{F}_q^n \to \mathbb{C}$ to be the additive character $\chi_b(x) = \omega_q^{\langle b,x\rangle}$ where $\omega_q = e^{2\pi i/q}$ is a primitive q-th root of unity. Note that this definition is valid even for $n = 1$.

Theorem 10.5. *Let f be an n-variate polynomial over $\mathbb{F}_q$ of degree d. Let f_d denote the degree-d homogeneous component of f and suppose that f_d is* smooth*, that is, the only common zero of the n polynomials $\{\partial f_d/\partial x_i\}_{i\in[n]}$ is the all-zero vector. Then for every non-zero $b \in \mathbb{F}_q$,*

$$\left|\sum_{x\in\mathbb{F}_q^n} \chi_b(f(x))\right| \leq (d-1)^n q^{n/2}.$$

Let $A \in \mathbb{F}^{m\times n}$ be a matrix where every m columns are linearly independent (e.g., a Vandermonde matrix). Let $d_1 < d_2 < \cdots < d_n$ be as in Proposition 10.3 and define the function $E : \mathbb{F}^n \to \mathbb{F}^m$ by

$$E(x_1, \ldots, x_n) = A \cdot \begin{pmatrix} x_1^{d_1} \\ \vdots \\ x_n^{d_n} \end{pmatrix}. \tag{28}$$

Theorem 10.6. *For every $0 < \beta < 1/2$, there exists a constant C such that the following holds: Let $k \leq n$ be integers and $\mathbb{F}$ be a prime field of size $q \geq n^{C\log\log n}$. Then for $m = \beta k$ the function $E : \mathbb{F}^n \to \mathbb{F}^m$ as in (28) is an affine extractor for min-entropy k with error $q^{-\Omega(k)}$. That is, for every affine subspace $V \subseteq \mathbb{F}^n$ of dimension k, if X_V is a random variable uniformly distributed on V, $E(X_V)$ is $q^{-\Omega(k)}$-close to uniform on $\mathbb{F}^k$.*

Proof. As we mentioned earlier, our proof is identical to the proof of Theorem 3.1 of [BDL16]. We provide the details for completeness.

Let $V \subseteq \mathbb{F}^n$ be an affine subspace of dimension k and let X be a random variable obtained by picking a random element from V and applying the extractor E above. Let ℓ be an affine map whose image is V as in Lemma 10.4.

By Lemma 3.5, it is enough to give an upper bound on $|\mathbb{E}[\chi_c(X)]|$ for every non-zero $c \in \mathbb{F}_q^n$.

Denote $b = c^T A$, and observe that

$$\chi_c(E(x)) = \omega_q^{b_1 x_1^{d_1} + \cdots + b_n x_n^{d_n}} = \chi_1(b_1 x_1^{d_1} + \cdots + b_n q x_n^{d_n}).$$

Hence,

$$|\,\mathbb{E}[\chi_c(X)]| = \left| q^{-k} \sum_{t \in \mathbb{F}_q^k} \chi_1(b_1 \ell_1(t)^{d_1} + \cdots + b_n \ell_n(t)^{d_n}) \right| . \tag{29}$$

Denote $D = \mathrm{LCM}(d_1, \ldots, d_n) \leq q^{\varepsilon}$ and $D_i = D/d_{j_i}$ for every $i \in [k]$. By performing the change of variables $t_i = s_i^{D_i}$ (which is invertible since $d_1, \ldots, d_n$ are all coprime to $q-1$), we define $\tilde{\ell}_j = \ell_j(s_1^{D_1}, \ldots, s_k^{D_k})$, so that (29) becomes

$$|\,\mathbb{E}[\chi_c(X)]| = \left| q^{-k} \sum_{s \in \mathbb{F}_q^k} \chi_1(b_1 \tilde{\ell}_1(s)^{d_1} + \cdots + b_n \tilde{\ell}_n(s)^{d_n}) \right| , \tag{30}$$

and the functions $\tilde{\ell}_1, \ldots, \tilde{\ell}_k$ have the following properties (as in Claim 3.5 in [Bou07]):

1. For every $i \in [k]$, $\tilde{\ell}_{j_i}^{d_{j_i}} = s_i^D$.

2. For all $j \notin \{j_1, \ldots, j_k\}$, $\tilde{\ell}_j^{d_j}$ is a polynomial in $s_1, \ldots, s_k$ of degree strictly less than D.

This implies that we can write (30) as

$$|\,\mathbb{E}[\chi_c(X)]| = \left| q^{-k} \sum_{s \in \mathbb{F}_q^k} \chi_1(b_{j_1} s_1^D + \cdots + b_{j_k} s_k^D + g(s)) \right| , \tag{31}$$

where g is a polynomial of degree strictly less than D.

Since c is non-zero and every m columns of A are linearly independent, the vector $b = c^T A$ has at most $m - 1 < k/2$ zero coordinates. Hence, at least $k/2$ of the values $b_{j_1}, \ldots, b_{j_k}$ are non-zero. Suppose without loss of generality that these are the first $k/2$ coordinates. Thus, we estimate (31) as

$$\begin{aligned} &|\,\mathbb{E}[\chi_c(X)]| \\ &\leq q^{-k/2} \sum_{s_{k/2+1}, \ldots, s_k \in \mathbb{F}_q} \left| q^{-k/2} \sum_{s_1, \ldots, s_{k/2} \in \mathbb{F}_q} \chi_1(b_{j_1} s_1^D + \cdots + b_{j_{k/2}} s_{k/2}^D + g_{s_{k/2+1}, \ldots, s_k}(s_1, \ldots, s_{k/2})) \right| \end{aligned}$$

with $b_{j_1}, \ldots, b_{j_{k/2}}$ non-zero and $\deg(g_{s_{k/2+1}, \ldots, s_k}) < d$ for every $s_{k/2+1}, \ldots, s_k$. That is, for every choice of $s_{k/2+1}, \ldots, s_k$, the degree D homogeneous component of the polynomial

$$b_{j_1} s_1^D + \cdots + b_{j_{k/2}} s_{k/2}^D + g_{s_{k/2+1}, \ldots, s_k}(s_1, \ldots, s_{k/2})$$

is smooth. By Theorem 10.5,

$$|\,\mathbb{E}[\chi_c(X)]| \le q^{-k/2} \cdot D^{k/2} \cdot q^{k/4} \le q^{(-1/4+\varepsilon/2)k}.$$

Setting $\varepsilon = 1/4 - \beta/2 > 0$ and applying Lemma 3.5, the statistical distance of X from the uniform distribution on $\mathbb{F}_q^m$ is at most

$$q^{(-1/4+\varepsilon/2)k} \cdot q^{m/2} \le q^{-(\varepsilon/2)k},$$

which concludes the proof. □

11 Explicit Noether Normalization for Affine Varieties and Affine Algebras

The Noether normalization lemma [Noe26, Nag62] is a cornerstone of commutative algebra and algebraic geometry. It states that any finitely generated commutative algebra over a field $\mathbb{F}$, or what we call an *affine algebra* over $\mathbb{F}$, is not too far from a polynomial ring, in the sense that it is always a finitely generated module over a subring that is isomorphic to a polynomial ring $\mathbb{F}[Y_1, \ldots, Y_k]$. The geometric interpretation of this statement is that any affine variety V over $\mathbb{F}$ is a "branched covering" of an affine space $\mathbb{A}_{\mathbb{F}}^k$, or more precisely, V admits a surjective finite morphism $\varphi_V : V \to \mathbb{A}_{\mathbb{F}}^k$.

When $\mathbb{F}$ is an infinite field (or more generally, a sufficiently large field), the polynomials that define the finite morphism φ_V may be chosen to be linear polynomials (see, e.g., Lemma 4.8). In general, φ_V can always be chosen to be defined by polynomials of sufficiently large degrees. In fact, counting arguments show that given the variety, a "random" polynomial map defined by polynomials of sufficiently large degrees would almost surely yield such a finite morphism. See [BE19] for a quantitative analysis. However, it is not known how to completely "derandomize" such counting arguments.

The first proof of the Noether normalization lemma for general affine algebras over arbitrary fields was given by Nagata [Nag53, Nag56, Nag62]. This proof has the interesting feature that it actually constructs a "universal" polynomial map $\varphi : \mathbb{A}_{\mathbb{F}}^n \to \mathbb{A}_{\mathbb{F}}^k$ that works for all low-degree affine varieties. Namely, for any low-degree affine variety $V \subseteq \mathbb{A}_{\mathbb{F}}^n$ of dimension k, the restriction of φ to V gives a finite morphism $\varphi|_V : V \to \mathbb{A}_{\mathbb{F}}^k$. The existence of such a polynomial map φ that is independent of V appears to be stronger and more intriguing than the existence of finite morphisms $V \to \mathbb{A}_{\mathbb{F}}^k$. In fact, we do not know how to prove the existence of φ via a counting argument.

While the polynomial map φ constructed by Nagata gives a uniform way of constructing finite morphisms, a drawback is that the degrees of the polynomials that define φ can get extremely high due to the iterative nature of the construction. More specifically, the map φ is constructed as a composition of polynomial maps $\varphi_i : \mathbb{A}_{\mathbb{F}}^{i+1} \to \mathbb{A}_{\mathbb{F}}^i$, $i = n-1, \ldots, k$ such that their restrictions $\varphi_i|_{V_{i+1}}$ are finite morphisms, where we inductively define $V_n = V$ and $V_i = \overline{\varphi_i(V_{i+1})}$ for $i = n-1, \ldots, k$. The problem is that composing with a polynomial map can increase the degree of a variety exponentially (see Lemma 4.12). The degree bound for the polynomials defining φ is at least doubly exponential for this reason.

Thus, it is a natural question to ask if there is a more efficient construction of the universal polynomial map φ. In this section, we show that the DKL construction in Section 6 is indeed such a construction, which always works when $|\mathbb{F}| \ge n$.

The construction of φ. We first recall the DKL construction in Section 6. Let $\mathbb{F}$ be a field. Let $n, d \in \mathbb{N}^+$ and $m, k \in [n]$. Let $d_1, \dots, d_n$ be n pairwise coprime integers greater than d. Let $M = (c_{i,j})_{i\in[m],j\in[n]} \in \mathbb{F}^{m\times n}$ be a k-regular matrix, i.e., any k distinct columns of M are linearly independent. Let $\varphi = \varphi(M) : \mathbb{A}^n_{\mathbb{F}} \to \mathbb{A}^m_{\mathbb{F}}$ be the polynomial map defined by $f_1, \dots, f_m \in \mathbb{F}[X_1, \dots, X_n]$, where $f_i := \sum_{j=1}^n c_{i,j} X_j^{d_j}$. In other words, φ is given by

$$\varphi : (a_1, \dots, a_n) \mapsto \left(\sum_{j=1}^n c_{1,j} a_j^{d_j}, \dots, \sum_{j=1}^n c_{m,j} a_j^{d_j} \right).$$

The main results of this subsection are the following theorems.

Theorem 11.1 (Explicit Noether normalization for affine varieties)**.** *Let V be an affine variety of dimension at most k and degree at most d over a field $\mathbb{F}$. Then $\varphi|_V : V \to \mathbb{A}^m_{\mathbb{F}}$ is a finite morphism.*

Theorem 11.1 translates into the following algebraic statement, Theorem 11.2, which gives an explicit Noether normalization lemma for affine algebras, i.e., finitely generated commutative algebras over a field.

Recall that the *Krull dimension* of a commutative ring A is the supremum of the lengths of all chains of prime ideals in A. If V is an affine variety over a field $\mathbb{F}$, then the Krull dimension of its coordinate ring $\mathbb{F}[V]$ is just the dimension of V.

Theorem 11.2 (Explicit Noether normalization for affine algebras)**.** *Suppose A is a commutative $\mathbb{F}$-algebra generated by $a_1, \dots, a_n \in A$ such that the Krull dimension of A is at most k. Let the ideal I of $\mathbb{F}[X_1, \dots, X_n]$ be the ideal of all polynomial relations satisfied by $a_1, \dots, a_n$. Also suppose the degree of the affine variety $V(I) \subseteq \mathbb{A}^n$ is at most d. Then A is a finitely generated module over its subring $S = \mathbb{F}[f_1(a), \dots, f_m(a)]$, where $f_1, \dots, f_m$ are the polynomials defining φ and $a = (a_1, \dots, a_n)$.*

The fact that A is a finitely generated module over S implies that the Krull dimension of S equals that of A. In the case where the Krull dimension of A is k and $k = m$, this means $f_1(a), \dots, f_m(a)$ are algebraically independent over $\mathbb{F}$, and hence S is isomorphic to a polynomial ring $\mathbb{F}[Y_1, \dots, Y_m]$ via $f_i(a) \mapsto Y_i$.

Theorem 11.1 and Theorem 11.2 are proved in Appendix D. The proof is inspired by and closely follows a geometric proof sketched in [KRS96, Remark 1].

Smaller fields. While $k \times n$ MDS matrices are generally not known over small finite fields $\mathbb{F}_q$, which prevents us from choosing $m = k$ over $\mathbb{F}_q$, it may still be possible to choose larger m for which (explicit) k-regular $m \times n$ matrices over $\mathbb{F}_q$ exist, and this would yield a finite morphism $\varphi|_V : V \to \mathbb{A}^m_{\mathbb{F}_q}$ by Theorem 11.1. As compositions of finite morphisms are finite [AM69, Corollary 5.4], by replacing n with m and V with $V' = \overline{\varphi(V)}$, we reduce the problem of constructing a finite morphism on $V \subseteq \mathbb{A}^n_{\mathbb{F}_q}$ to constructing that on $V' \subseteq \mathbb{A}^m_{\mathbb{F}_q}$, where V' has the same dimension as V but lives in a possibly much smaller affine space $\mathbb{A}^m_{\mathbb{F}_q}$. The degree of V', however, may be significantly larger than that of V. See Lemma 4.12 for a general upper bound on the degree.

For example, while we do not know the existence of $k \times n$ MDS matrices over small finite fields $\mathbb{F}_q$, one can still use a BCH-code-like construction to obtain an $m \times n$ k-regular matrix with

$m = O(k \log_q n)$, which can be much smaller than n if $k \ll n$. Applying the resulting map φ reduces the dimension of the ambient space from n to m.

However, when q is really small and k is close to n, it may be possible that one can only choose $m = n-1$ and hence only reduce the dimension of the ambient space by one at each step. This is essentially the same method used in Nagata's construction. Currently, all constructions of the universal polynomial map $\varphi : \mathbb{A}^n_{\mathbb{F}_q} \to \mathbb{A}^k_{\mathbb{F}_q}$ with $k = \dim V$ that we know over a constant-size field $\mathbb{F}_q$ use polynomials of degree at least doubly exponential in $\min\{k, n-k\}$ due to the blow-up of the degree of the variety. It is an interesting question to ask if there exist constructions with a better degree bound over constant-size fields.

Appendices

A Noether Normalization via Linear Maps

In this section, we prove the quantitative Noether normalization lemma (Lemma 4.8) and its variant, Lemma 4.9. We also explain how to remove the condition $q > 2(k+1)d^2$ in [CM06, Theorem 7.1] and obtain Theorem 4.6.

Although Lemma 4.8 concerns only affine varieties, we need to deal with projective varieties that "complete" affine varieties.

Projective spaces and projective varieties. Fix $\mathbb{F}$ to be an algebraically closed field. The *projective n-space* $\mathbb{P}^n$ over $\mathbb{F}$, as a set, is the quotient set $\left(\mathbb{F}^{n+1} \setminus \{\mathbf{0}\}\right) / \sim$, where $\mathbf{0}$ is the origin $(0, \dots, 0)$ and $\sim$ is the equivalence relation defined by scaling, i.e., $u \sim v$ if $u = cv$ for some $c \in \mathbb{F}^\times$. We use $(n+1)$-tuples $(x_0, \dots, x_n)$ to represent points in $\mathbb{P}^n$.

We equip $\mathbb{P}^n$ with the *Zariski topology* over $\mathbb{F}$, where a subset is closed if it is the set of common zeros of a set of homogeneous polynomials in $\mathbb{F}[X_0, X_1, \dots, X_n]$. Call $X_0, \dots, X_n$ the *homogeneous coordinates* of $\mathbb{P}^n$. A closed subset $V \subseteq \mathbb{P}^n$ is said to be a *projective variety* over $\mathbb{F}$.

The projective space $\mathbb{P}^n$ is covered by open subsets $U_i := \{(x_0, \dots, x_n) \in \mathbb{P}^n : x_i \neq 0\}$, $i = 0, \dots, n$. Each U_i can be identified with the affine space $\mathbb{A}^n$ via the map

$$(x_0, \dots, x_n) \mapsto \left(\frac{x_0}{x_i}, \dots, \frac{x_{i-1}}{x_i}, \frac{x_{i+1}}{x_i}, \dots, \frac{x_n}{x_i}\right).$$

And the subspace topology of U_i induced from that of $\mathbb{P}^n$ is precisely the Zariski topology of $\mathbb{A}^n$ if we identify U_i with $\mathbb{A}^n$ this way.

While the sets U_i are symmetric, we often fix $i = 0$ and regard $\mathbb{A}^n$ as an open subset of $\mathbb{P}^n$ by identifying it with $U_0 \subseteq \mathbb{P}^n$ as above. Its complement $\mathbb{P}^n \setminus \mathbb{A}^n$ may be identified with $\mathbb{P}^{n-1}$ (for $n > 0$) via the map $(x_0, \dots, x_n) \mapsto (x_1, \dots, x_n)$, and is called the *hyperplane at infinity.* The *projective closure* of an affine variety $V \subseteq \mathbb{A}^n$, which we denote by V_{cl}, is the smallest projective subvariety of $\mathbb{P}^n$ that contains V as a subset. And we have $V_{\mathrm{cl}} \cap \mathbb{A}^n = V$.

The notions of irreducibility, irreducible components, degree, and dimension all extend to projective varieties. We have $\deg V_{\mathrm{cl}} = \deg V$ and $\dim V_{\mathrm{cl}} = \dim V$ for an affine variety $V \subseteq \mathbb{A}^n$, i.e., taking the projective closure preserves the degree and the dimension.

The notions of morphisms and finite morphisms also extend to projective varieties. See, e.g., [Sha94]. If $\varphi : V \to V'$ is a finite morphism between projective varieties, and $\varphi|_{\varphi^{-1}(U)} : \varphi^{-1}(U) \to U$

is a morphism between affine varieties for some open subset U of V', then $\varphi|_{\varphi^{-1}(U)}$ is also a finite morphism ([Sha94, §I.5.3, Theorem 5]).

The following lemma gives a way of finding a finite morphism $V \to \mathbb{P}^k$ defined by linear polynomials.

Lemma A.1 ([Sha94, §I.5.3, Theorem 7])**.** *Let $V \subseteq \mathbb{P}^n$ be a projective variety over $\mathbb{F}$. Suppose $\ell_0, \dots, \ell_k \in \mathbb{F}[X_0, \dots, X_n]$ are (homogeneous) linear polynomials that have no common zero on V. Then they define a finite morphism $V \to \mathbb{P}^k$ that sends $a \in V$ to $(\ell_0(a), \dots, \ell_k(a))$.*

Thus, the problem of finding a finite morphism from a projective variety V to $\mathbb{P}^k$ reduces to finding $k+1$ linear polynomials that have no common zero on V. We use *Chow forms* to show the existence of such linear polynomials.

Lemma A.2. *Let $V \subseteq \mathbb{P}^n$ be an irreducible projective variety over $\mathbb{F}$ of dimension k and degree d. There exists a nonzero polynomial $R_V \in \mathbb{F}[Y_{0,0}, \dots, Y_{k,n}]$ such that for all $c_{0,0}, \dots, c_{k,n} \in \mathbb{F}$, the linear polynomials $\ell_0, \dots, \ell_k$ defined by $\ell_i := \sum_{j=0}^n c_{i,j} X_j$ have no common zero on V iff $R_V(c_{0,0}, \dots, c_{k,n}) \neq 0$. Moreover, R_V is multihomogeneous of degree $(d, \dots, d)$ in*

$$\{Y_{0,0}, \dots, Y_{0,n}\}, \dots, \{Y_{k,0}, \dots, Y_{k,n}\},$$

i.e., it is homogeneous of degree d in each of the $k+1$ group of variables.

The polynomial R_V is unique up to a scalar and is called the *Chow form* of V. See, e.g., [KPS01].

We also need the following lemma on the existence of a non-root of a polynomial.

Lemma A.3. *Suppose $P \in \mathbb{F}[X_{1,1}, \dots, X_{k,n}]$ is a polynomial such that for all $i \in [k]$, the (total) degree of P in the variables $X_{i,1}, \dots, X_{i,n}$ is at most d. Let S be a finite subset of $\mathbb{F}$ of size greater than d. Then P has a non-root in S^{kn}.*

Proof. We assign the values $c_{i,1}, \dots, c_{i,n} \in S$ to the k groups of variables $\{X_{i,1}, \dots, X_{i,n}\}$ for $i = 1, \dots, k$ one by one, such that at each step, the polynomial P remain nonzero. This is obviously true at the beginning. Now at the beginning of the i-th step, assume that P remains nonzero after the partial assignment $X_{i',j} = c_{i',j}$ for $i' < i$ and $j \in [n]$, and call this polynomial P_i. View P_i as a polynomial in the variables $X_{i+1,1}, \dots, X_{k,n}$ over the ring $\mathbb{F}[X_{i,1}, \dots, X_{i,n}]$. Then as $P_i \neq 0$, it has a term whose coefficient Q is a nonzero polynomial of degree at most d in $X_{i,1}, \dots, X_{i,n}$. By the DeMillo–Lipton–Schwartz–Zippel lemma [Sch80, Zip79, DL78], there exist $c_{i,1}, \dots, c_{i,n} \in S$ such that $Q(c_{i,1}, \dots, c_{i,n}) \neq 0$, and hence P_i remain nonzero after the assignment $X_{i,j} = c_{i,j}$, $j = 1, \dots, n$. Continuing this process, we see that P has a non-root in S^{kn}. □

We are now ready to prove Lemma 4.8. For convenience, we first restate this lemma.

Lemma 4.8 (Noether normalization)**.** *Let $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ be an affine variety of dimension k and degree d over a field $\mathbb{F}$. Suppose S is a finite subset of $\mathbb{F}$ of size greater than d. Then there exists a polynomial map $\varphi : \mathbb{A}^n_{\mathbb{F}} \to \mathbb{A}^k_{\mathbb{F}}$ defined by linear polynomials $\ell_i = \sum_{j=1}^n c_{i,j} X_i \in \mathbb{F}[X_1, \dots, X_n]$ with coefficients $c_{i,1}, \dots, c_{i,n} \in S$ for $i = 1, \dots, k$ such that $\varphi|_V : V \to \mathbb{A}^k_{\mathbb{F}}$ is a finite morphism.*

Proof. We may assume that $\mathbb{F}$ is algebraically closed. This is because finiteness of morphisms over $\mathbb{F}$ follows from that over $\overline{\mathbb{F}}$ by a descent argument (see Appendix D).

Suppose $V_1, \dots, V_s$ are the irreducible components of V, where $\dim V_t = k_t \le k$ and $\deg V_t = d_t$ for $t \in [s]$. Identify $\mathbb{A}^n$ with the open subset $U_0 = \{(x_0, \dots, x_n) \in \mathbb{P}^n : x_0 \ne 0\}$ of $\mathbb{P}^n$, and let $H = \mathbb{P}^n \setminus U_0$ be the hyperplane at infinity.

For $t \in [s]$, let

$$P_t = R_{(V_t)_{\mathrm{cl}}} \in \mathbb{F}[Y_{0,0}, \dots, Y_{k_t,n}] \subseteq \mathbb{F}[Y_{0,0}, \dots, Y_{k,n}].$$

That is, P_t is the Chow form of the projective closure of V_t. Then P_t is multihomogeneous of degree $(d_t, \dots, d_t)$. For $t \in [s]$, let $\widehat{P}_t \in \mathbb{F}[Y_{1,1}, \dots, Y_{n,n}]$ be the polynomial obtained from P_t by assigning $(Y_{0,0}, Y_{0,1}, \dots, Y_{0,n}) = (1, 0, \dots, 0)$ and $Y_{1,0} = \cdots = Y_{k,0} = 0$. Then each $\widehat{P}_t$ remains a nonzero polynomial. In fact, $\widehat{P}_t$ may be viewed as the Chow form of $(V_t)_{\mathrm{cl}} \cap H$ as a projective subvariety of $H \cong \mathbb{P}^{n-1}$ of dimension $k-1$. (If $\dim V_t = 0$, then $(V_t)_{\mathrm{cl}} \cap H = \emptyset$ and $\widehat{P}_t$ is a nonzero constant.)

Suppose $c_{1,1}, \dots, c_{k,n} \in S$ satisfy $\widehat{P}_t(c_{0,0}, \dots, c_{n,k}) \ne 0$ for all $t \in [s]$. We claim the linear forms $\ell_i = \sum_{j=1}^n c_{i,j} X_i$ define a finite morphism $V \to \mathbb{A}^k$. To see this, note that by the property of Chow forms, the polynomials $\ell_0 := X_0$ and $\ell_1, \dots, \ell_k$ have no common zero on $V_{\mathrm{cl}} = \bigcup_{t=1}^s (V_t)_{\mathrm{cl}}$. So by Lemma A.1, $\ell_0, \dots, \ell_k$ define a finite morphism $\varphi : V_{\mathrm{cl}} \to \mathbb{P}^k$. Let U_0' be the open subset $\{(x_0, \dots, x_k) \in \mathbb{P}^k : x_0 \ne 0\}$ of $\mathbb{P}^k$. As $\ell_0 = X_0$, we have $\ell_0(x) \ne 0$ for $x \notin U_0$ and hence $\varphi(H) \cap U_0' = \emptyset$. So $\varphi^{-1}(U_0') = U_0 \cap V_{\mathrm{cl}} = V$. As φ is a finite morphism, so is $\varphi|_V : V \to U_0'$. Finally, note that $\varphi|_V$ is exactly the morphism $V \to \mathbb{A}^k$ defined by $\ell_1, \dots, \ell_k$ if we identify U_0 with $\mathbb{A}^k$.

So it remains to prove the existence of $c_{1,1}, \dots, c_{k,n} \in S$ such that $\widehat{P}_t(c_{1,1}, \dots, c_{k,n}) \ne 0$ for $t \in [s]$. This follows by applying Lemma A.3 to $P := \prod_{t=1}^s \widehat{P}_t$ and noting that P is multihomogeneous of degree $(d, \dots, d)$. □

One can prove an analogue of Lemma 4.8 for projective varieties with a similar, and in fact, simpler proof. But we only need Lemma 4.8 for affine varieties in this paper.

Proof of Lemma 4.9. The proof is the same as that of Lemma 4.8, except that we consider V_1 and V_2 simultaneously and apply the union bound when picking each linear polynomial ℓ_i. □

The effective Lang–Weil bound. Lemma A.3 above can also be used to prove the effective Lang–Weil bound (Theorem 4.6). We restate the theorem below for convenience.

Theorem 4.6 (Effective Lang–Weil bound). *Let $V \subseteq \mathbb{A}^n_{\mathbb{F}_q}$ be an absolutely irreducible affine variety over $\mathbb{F}_q$ of dimension k and degree d. Then*

$$|V(\mathbb{F}_q) - q^k| < (d-1)(d-2)q^{k-1/2} + 5d^{13/3}q^{k-1}.$$

In particular, we have $|V(\mathbb{F}_q)| \ge q^k/2$ if $q \ge 20d^5$.

This bound was proved as [CM06, Theorem 7.1] with an extra condition that $q > 2(k+1)d^2$. We first explain why this condition was assumed in [CM06].

To prove [CM06, Theorem 7.1], Cafure and Matera first proved an effective Lang–Weil bound when V is an absolutely irreducible *hypersurface* in $\mathbb{A}^n_{\mathbb{F}_q}$ without the condition $q > 2(k+1)d^2$ (see [CM06, Theorem 5.2]). To extend it to the general case, they further argued that if $q > 2(k+1)d^2$, then there exists an affine linear map $\pi : \mathbb{A}^n_{\mathbb{F}_q} \to \mathbb{A}^{k+1}_{\mathbb{F}_q}$ over $\mathbb{F}_q$, defined by polynomials $\sum_{j=1}^n \lambda_{ij} X_j + \gamma_i \in \mathbb{F}_q[X_1, \dots, X_n]$, $i = 1, \dots, k+1$, that induces a *birational equivalence* $\pi|_V : V \dashrightarrow H$ between V and a hypersurface $H \subseteq \mathbb{A}^{k+1}_{\mathbb{F}_q}$. Such a birational equivalence may be thought of as an "almost isomorphism" between V and H. In particular, $\pi|_V$ is an "almost bijection" between the set of rational points of V and that of H, i.e., there exist dense open subsets $U \subseteq V$ and $U' \subseteq H$ such that the rational points in U are mapped bijectively to those in U' under $\pi|_V$. Furthermore, the sets $(V \setminus U)(\mathbb{F}_q)$ and $(H \setminus U')(\mathbb{F}_q)$ are small. The effective Lang–Weil bound for affine varieties then easily follows from the bound for hypersurfaces.

To see intuitively why such a morphism π should exist, note that the quantitative Noether normalization lemma (Lemma 4.8) already guarantees the existence of a morphism $V \to \mathbb{A}^k_{\mathbb{F}_q}$ that is finite (and hence surjective). However, this map is finite-to-one instead of (almost) one-to-one. But note that a morphism $\mathbb{A}^n_{\mathbb{F}_q} \to \mathbb{A}^{k+1}_{\mathbb{F}_q}$ has one extra coordinate in the output, and we can use this extra coordinate to distinguish the finitely many preimages of a general point $b \in \mathbb{A}^k_{\mathbb{F}_q}$. Thus, we should expect that a general affine linear map $\pi : \mathbb{A}^n_{\mathbb{F}_q} \to \mathbb{A}^{k+1}_{\mathbb{F}_q}$ defines an "almost isomorphism" from V to a hypersurface. Indeed, Cafure and Matera used the *Chow form* and the *discriminant* to prove the existence of such a morphism π. Specifically, they showed that there exists a nonzero polynomial G in the variables $(\Lambda_{i,j})_{i\in[k+1],j\in[n]}$ and $(\Gamma_i)_{i\in[k+1]}$ over $\overline{\mathbb{F}}_q$ such that π is a desired morphism whenever the coefficients $(\lambda_{i,j})_{i\in[k+1],j\in[n]}$ and $(\gamma_i)_{i\in[k+1]}$ of the polynomials that define π form a non-root of G [CM06, Theorem 6.1]. Moreover, G has the property that its degree in $\Lambda_{i,1}, \dots, \Lambda_{i,n}, \Gamma_i$ is at most $2d^2$ for $i \in [k+1]$ (see the proof of [CM06, Theorem 6.1]). In particular, its total degree is at most $2(k+1)d^2$.

Next, Cafure and Matera used essentially the DeMillo–Lipton–Schwartz–Zippel lemma [Sch80, Zip79, DL78] and the fact that $\deg G \le 2(k+1)d^2$ to argue that a non-root of G exists in $\mathbb{F}_q^{(k+1)(n+1)}$, and this is where they need the condition $q > 2(k+1)d^2$ [CM06, Corollary 6.2].

However, we can take advantage of the fact that the degree of G in each group of variables $\Lambda_{i,1}, \dots, \Lambda_{i,n}, \Gamma_i$ is at most $2d^2$ for $i \in [k+1]$, and use Lemma A.3 instead. This immediately allows us to relax the condition $q > 2(k+1)d^2$ to $q > 2d^2$.

Finally, observe that when $q \le 2d^2$, the bound $|V(\mathbb{F}_q) - q^k| < (d-1)(d-2)q^{k-1/2} + 5d^{13/3}q^{k-1}$ follows from the trivial lower bound $|V(\mathbb{F}_q)| \ge 0$ and the upper bound $|V(\mathbb{F}_q)| \le dq^k$ in Lemma 4.5. So we can remove the condition about q completely.

B The Effective Fiber Dimension Theorem

We prove the general form of the effective fiber dimension theorem (Theorem 4.10) in this section. The base field $\mathbb{F}$ is assumed to be an algebraically closed field.

Generalized Perron Theorem. We need the following result proved by Jelonek [Jel05], generalizing a classical result of Perron [Per27] that bounds the degree of annihilating polynomials for algebraic dependent polynomials.

Theorem B.1 (Generalized Perron Theorem [Jel05, Theorem 3.3])**.** *Suppose $f : \mathbb{A}^n \to \mathbb{A}^m$ is a polynomial map defined by polynomials $f_1, \dots, f_m \in \mathbb{F}[X_1, \dots, X_n]$, where $\deg f_i = d_i > 0$. Let $V \subseteq \mathbb{A}^n$ be an irreducible affine variety over $\mathbb{F}$, and let $W = \overline{f(V)} \subseteq \mathbb{A}^m$. Suppose $\dim V = \dim W = m-1$. Then there exists a nonzero polynomial $Q \in \mathbb{F}[Y_1, \dots, Y_m]$ such that $Q(f_1, \dots, f_m)$ vanishes identically on V and $\deg(Q(Y_1^{d_1}, \dots, Y_m^{d_m})) \le \deg V \cdot \prod_{i=1}^m d_i$.*

Proof of the effective fiber dimension theorem. Now we are ready to prove Theorem 4.10. For convenience, we first restate the theorem.

Theorem 4.10 (Effective fiber dimension theorem – general form)**.** *Let $V \subseteq \mathbb{A}^n$ be an irreducible affine variety of dimension k over an algebraically closed field $\mathbb{F}$. Let $h_1, \dots, h_s \in \mathbb{F}[X_1, \dots, X_n]$ with* $\deg h_1 \geq \cdots \geq \deg h_s$. *Let $f_1, \dots, f_m \in \mathcal{L}_{h_1,\dots,h_s,\mathbb{F}}$, which define a polynomial map $f : \mathbb{A}^n \to \mathbb{A}^m$. Let $k' = \dim \overline{f(V)}$.*

Let $j_1, \dots, j_{k'} \in [m]$ such that the morphism $f' : V \to \mathbb{A}^{k'}$ defined by $f_{j_1}, \dots, f_{j_{k'}}$ is dominant, which exist by Lemma 4.3. Let $V_{f'} \subseteq \mathbb{A}^n_{\mathbb{F}(Y_1,\dots,Y_{k'})}$ be the generic fiber of f' (see the definition after Theorem 4.4). Finally, let $\ell_1, \dots, \ell_k \in \mathbb{F}[X_1, \dots, X_n]$ be linear polynomials such that both the morphism $\pi : V \to \mathbb{A}^k$ defined by $\ell_1, \dots, \ell_k$ and the morphism $\tau : V_{f'} \to \mathbb{A}^{k-k'}_{\mathbb{F}(Y_1,\dots,Y_{k'})}$ defined by $\ell_1, \dots, \ell_{k-k'}$ are finite.

Let $t \in \{0, \dots, k - k'\}$. Then there exists a polynomial $P \in \mathbb{F}[X_1, \dots, X_n]$ of degree at most $k' \cdot \deg V \cdot \prod_{i=1}^{k'} \deg h_i$ *that does not vanish identically on V such that the following holds: Let $\varphi : \mathbb{A}^n \to \mathbb{A}^{t+m}$ be the polynomial map defined by $\ell_1, \dots, \ell_t, f_1, \dots, f_m$. Then for every $a \in V$ satisfying $P(a) \neq 0$, the fiber $\varphi|_V^{-1}(\varphi(a))$ is equidimensional of dimension $k - k' - t$.*

One special form of the effective fiber dimension theorem with $V = W = \mathbb{A}^k$ was essentially proved in [GSS19] using Perron's bound [Per27]. Our proof of Theorem 4.10 can be seen as a generalization of this proof. It can also be seen as an effective version of a standard proof of the fiber dimension theorem (see [Vak22, Proof of Theorem 12.4.1]).

Towards proving Theorem 4.10, we first prove the following lemma using the Generalized Perron Theorem.

Lemma B.2. *Use the notations in Theorem 4.10 and assume that* $\deg h_i > 0$ *for $i \in [s]$. Suppose* $\dim V = \dim \overline{f(V)} = k = m$. *Also suppose at least u polynomials among $f_1, \dots, f_k$ are linear. Let $g \in \mathbb{F}[X_1, \dots, X_n]$ such that* $\deg g > 0$. *Then there exists a nonzero polynomial $Q \in \mathbb{F}[Y_1, \dots, Y_{k+1}]$ satisfying the following:*

1. *$Q(f_1, \dots, f_k, g) \in \mathbb{F}[X_1, \dots, X_n]$ vanishes identically on V.*

2. *View Q as a univariate polynomial in Y_{k+1} over $\mathbb{F}[Y_1, \dots, Y_k]$ and let $Q^* \in \mathbb{F}[Y_1, \dots, Y_k]$ be its leading coefficient. Then the degree of $Q^*(f_1, \dots, f_k) \in \mathbb{F}[X_1, \dots, X_n]$ is at most* $\deg V \cdot \deg g \cdot \prod_{i=1}^{\min\{k-u,s\}} \deg h_i$.

Proof. Note that changing the coordinate system of $\mathbb{A}^m = \mathbb{A}^k$ via an invertible linear transformation does not affect the statement. So by permuting the polynomials f_i, we may assume the linear polynomials f_i appear at the end of the list $f_1, \dots, f_k$. By applying Gaussian elimination, we may further assume $f_i \in \mathcal{L}_{h_i,\dots,h_s,\mathbb{F}}$ (i.e., f_i does not depend on $h_1, \dots, h_{i-1}$) for $i = 1, \dots, \min\{k-u, s\}$, and $\deg(f_i) \leq 1$ for $i = \min\{k-u, s\}+1, \dots, k$. In particular, we have $\deg f_i \leq \deg h_i$ for $i = 1, \dots, \min\{k-u, s\}$.

Let ψ be the polynomial map $\mathbb{A}^n \to \mathbb{A}^{k+1}$ defined by $f_1, \dots, f_k, g$. The dimension of $\overline{\psi(V)}$ is k since it cannot exceed $k = \dim V$ and the dimension of $\overline{f(V)}$ is already k. As $\dim \overline{f(V)} = k$, we necessarily have $\deg f_i > 0$ for $i \in [k]$. Applying the Generalized Perron Theorem (Theorem B.1) to ψ, we see that there exists a nonzero polynomial $Q \in \mathbb{F}[Y_1, \dots, Y_{k+1}]$ such that $Q(f_1, \dots, f_k, g)$ vanishes identically on V and

$$\deg\left(Q\left(Y_1^{\deg f_1}, \dots, Y_k^{\deg f_k}, Y_{k+1}^{\deg g}\right)\right) \leq \deg V \cdot \deg g \cdot \prod_{i=1}^{k} \deg f_i \leq \deg V \cdot \deg g \cdot \prod_{i=1}^{\min\{k-u,s\}} \deg h_i.$$

Let $Q^* \in \mathbb{F}[Y_1, \dots, Y_k]$ be the leading term of Q in Y_{k+1}. Then

$$\begin{aligned}\deg(Q^*(f_1, \dots, f_k)) &\leq \deg\left(Q^*\left(Y_1^{\deg f_1}, \dots, Y_k^{\deg f_k}\right)\right)\\ &\leq \deg\left(Q\left(Y_1^{\deg f_1}, \dots, Y_k^{\deg f_k}, Y_{k+1}^{\deg g}\right)\right)\\ &\leq \deg V \cdot \deg g \cdot \prod_{i=1}^{\min\{k-u,s\}} \deg h_i.\end{aligned}$$

To see that the first inequality holds, note that the monomials of $Q^*\left(Y_1^{\deg f_1}, \dots, Y_k^{\deg f_k}\right)$ correspond one-to-one to the monomials of $Q^*(Y_1, \dots, Y_k)$, and substituting f_i for $Y_i^{\deg f_i}$ within each monomial does not increase its degree. □

We use Lemma B.2 to prove Theorem 4.10.

Proof of Theorem 4.10. We may assume $\deg h_i > 0$ for $i \in [s]$ by removing all polynomials h_i that are constants. Let $\psi : \mathbb{A}^n \to \mathbb{A}^k$ be the polynomial map defined by $\ell_1, \dots, \ell_{k-k'}, f_{j_1}, \dots, f_{j_{k'}}$. As τ is finite, we have $\dim \overline{\psi(V)} = \dim V = k$. And we have $\dim \overline{\varphi(V)} = t + k'$ as otherwise the dimension of $\overline{\psi(V)}$ cannot achieve k.

Consider $i \in \{k - k' + 1, \dots, k\}$. Applying Lemma B.2 to $\psi : \mathbb{A}^n \to \mathbb{A}^k$ (which is defined by $\ell_1, \dots, \ell_{k-k'}, f_{j_1}, \dots, f_{j_{k'}} \in \mathcal{L}_{h_1,\dots,h_s,\ell_1,\dots,\ell_{k-k'},\mathbb{F}}$) with $g = \ell_i$ and $u = k - k'$, we see that there exists a nonzero polynomial $Q_i \in \mathbb{F}[Y_1, \dots, Y_{k+1}]$ satisfying the following:

1. $Q_i(\ell_1, \dots, \ell_{k-k'}, f_{j_1}, \dots, f_{j_{k'}}, \ell_i)$ vanishes identically on V.
2. Let $Q_i^* \in \mathbb{F}[Y_1, \dots, Y_k]$ be the leading coefficient of Q_i in Y_{k+1}. Then the degree of $Q_i^*(\ell_1, \dots, \ell_{k-k'}, f_{j_1}, \dots, f_{j_{k'}}) \in \mathbb{F}[X_1, \dots, X_n]$ is at most $\deg V \cdot \prod_{i=1}^{k'} \deg h_i$.

For $g \in \mathbb{F}[X_1, \dots, X_n]$, let $\bar{g} := g + I(V) \in \mathbb{F}[V]$. Then $\bar{\ell}_1, \dots, \bar{\ell}_{k-k'}, \bar{f}_{j_1}, \dots, \bar{f}_{j_{k'}}$ are algebraically independent over $\mathbb{F}$ as $\psi|_V : V \to \mathbb{A}^k$ is dominant. In particular, we have $Q_i^*(\bar{\ell}_1, \dots, \bar{\ell}_{k-k'}, \bar{f}_{j_1}, \dots, \bar{f}_{j_{k'}}) \neq 0$, or equivalently, $Q_i^*(\ell_1, \dots, \ell_{k-k'}, f_{j_1}, \dots, f_{j_{k'}})$ does not vanish on V.

The Zariski closure of the image of V under the polynomial map $\mathbb{A}^n \to \mathbb{A}^{k+1}$ defined by $\ell_1, \dots, \ell_{k-k'}, f_{j_1}, \dots, f_{j_{k'}}, \ell_i$ is a hypersurface of $\mathbb{A}^{k+1}$ defined by a single polynomial, which we may assume to be Q_i. So Q_i generates the ideal of the polynomial relations satisfied by $\bar{\ell}_1, \dots, \bar{\ell}_{k-k'}, \bar{f}_{j_1}, \dots, \bar{f}_{j_{k'}}, \bar{\ell}_i$.

Let $\mathbb{K} := \mathbb{F}(\bar{f}_{j_1}, \dots, \bar{f}_{j_{k'}})$, which can be viewed as the function field of the target $\mathbb{A}^{k'}$ of the dominant morphism f'. By the finiteness of τ, the coordinate ring $\mathbb{K}[V_{f'}]$ of $V_{f'}$ is a finitely generated module over $\mathbb{K}[\bar{\ell}_1, \dots, \bar{\ell}_{k-k'}]$. So $\bar{\ell}_i$ is integral over $\mathbb{K}[\bar{\ell}_1, \dots, \bar{\ell}_{k-k'}]$, i.e., it is a root of a monic polynomial over the ring $\mathbb{K}[\bar{\ell}_1, \dots, \bar{\ell}_{k-k'}]$ [AM69, Proposition 5.1]. It follows that the leading coefficient Q_i^* of Q_i does not depend on $Y_1, \dots, Y_{k-k'}$, although it may depend on $Y_{k-k'+1}, \dots, Y_k$ since $\bar{f}_{j_1}, \dots, \bar{f}_{j_{k'}}$ are invertible in $\mathbb{K}$. So $Q_i^* \in \mathbb{F}[Y_{k-k'+1}, \dots, Y_k]$.

Let $P = \prod_{i=k-k'+1}^{k} Q_i^*(f_{j_1}, \dots, f_{j_{k'}})$. Then P does not vanish identically on V and has degree at most $k' \cdot \deg V \cdot \prod_{i=1}^{k'} \deg h_i$.

Consider $a \in V$ such that $P(a) \neq 0$, and let Z be an irreducible component of $\varphi|_V^{-1}(\varphi(a))$. Note that $\dim Z \geq \dim V - \dim \overline{\varphi(V)} = \dim V - (t + k') = k - k' - t$ by the first claim in the fiber dimension theorem (Theorem 4.4).

It remains to prove that $\dim Z \leq k - k' - t$. For $g \in \mathbb{F}[X_1, \dots, X_n]$, let $\bar{\bar{g}} := g + I(Z) \in \mathbb{F}[Z]$. Note that $\bar{\bar{\ell}}_i = \ell_i(a) \in \mathbb{F}$ for $i \in [t]$ since $\ell_i - \ell_i(a)$ vanishes identically on $\varphi|_V^{-1}(\varphi(a)) \supseteq Z$. Similarly, we have $\bar{\bar{f}}_{j_i} = f_{j_i}(a) \in \mathbb{F}$ for $i \in [k']$ since $f_{j_i} - f_{j_i}(a)$ vanishes identically on $\varphi|_V^{-1}(\varphi(a)) \supseteq Z$.

For $i \in \{k - k' + 1, \dots, k\}$, as $Q_i(\ell_1, \dots, \ell_{k-k'}, f_{j_1}, \dots, f_{j_{k'}}, \ell_i)$ vanishes identically on $V \supseteq Z$, we see that $\bar{\bar{\ell}}_i$ is a root of the univariate polynomial

$$\begin{aligned} & Q_i(\bar{\bar{\ell}}_1, \dots, \bar{\bar{\ell}}_{k-k'}, \bar{\bar{f}}_{j_1}, \dots, \bar{\bar{f}}_{j_{k'}}, Y_{k+1}) \\ =& Q_i(\ell_1(a), \dots, \ell_t(a), \bar{\bar{\ell}}_{t+1}, \dots, \bar{\bar{\ell}}_{k-k'}, f_{j_1}(a), \dots, f_{j_{k'}}(a), Y_{k+1}) \in (\mathbb{F}[\bar{\bar{\ell}}_{t+1}, \dots, \bar{\bar{\ell}}_{k-k'}])[Y_{k+1}]. \end{aligned}$$

Its leading coefficient is $Q_i^*(f_{j_1}(a), \dots, f_{j_{k'}}(a))$, which is a nonzero element in $\mathbb{F}$ since $P(a) \neq 0$. So $\bar{\bar{\ell}}_i$ is a root of a monic polynomial over $\mathbb{F}[\bar{\bar{\ell}}_{t+1}, \dots, \bar{\bar{\ell}}_{k-k'}]$ for $i \in \{k - k' + 1, \dots, k\}$. It follows that $\mathbb{F}(\bar{\bar{\ell}}_1, \dots, \bar{\bar{\ell}}_k)$ is a finite extension of $\mathbb{F}(\bar{\bar{\ell}}_{t+1}, \dots, \bar{\bar{\ell}}_{k-k'})$. On the other hand, as the morphism $\pi : V \to \mathbb{A}^k$ defined by $\ell_1, \dots, \ell_k$ is finite and Z is an affine subvariety of V, we know $\mathbb{F}[Z]$ is a finitely generated module over $\mathbb{F}[\bar{\bar{\ell}}_1, \dots, \bar{\bar{\ell}}_k]$. So $\mathbb{F}(Z)$ is a finite extension of $\mathbb{F}(\bar{\bar{\ell}}_1, \dots, \bar{\bar{\ell}}_k)$, and hence also a finite extension of $\mathbb{F}(\bar{\bar{\ell}}_{t+1}, \dots, \bar{\bar{\ell}}_{k-k'})$. Therefore, the trancendence degree of $\mathbb{F}(Z)$ over $\mathbb{F}$ is at most $k - k' - t$, i.e., $\dim Z \leq k - k' - t$. □

C Miscellanea

Absolute irreducibility. The following lemma gives alternative characterizations of absolute irreducibility when the finite field $\mathbb{F}_q$ is large enough.

Lemma C.1. *Let V be a nonempty affine variety over $\mathbb{F}_q$ of dimension k and degree d. Suppose $q \geq 20d^5$. Then the following are equivalent:*

(1) At least one irreducible component of V of dimension k is absolutely irreducible.

(2) $|V(\mathbb{F}_q)| \geq q^k/2$.

(3) $|V(\mathbb{F}_q)| > d^2 q^{k-1}$.

Proof. The fact that (1) implies (2) follows from the effective Lang–Weil bound (Theorem 4.6). And (2) implies (3) as $q \geq 20d^5$. Finally, to see that (3) implies (1), suppose that none of the irreducible components of V of dimension k is absolutely irreducible. Let V' be the union of the irreducible components of V of dimension k, and let V'' be the union of the remaining irreducible components. Then we have $|V'(\mathbb{F}_q)| \leq (\deg V')^2 q^{k-1}$ by Lemma 7.2 and $|V''(\mathbb{F}_q)| \leq \deg V'' q^{k-1}$ by Lemma 4.5. It follows that

$$|V(\mathbb{F}_q)| \leq |V'(\mathbb{F}_q)| + |V''(\mathbb{F}_q)| \leq (\deg V' + \deg V'')^2 q^{k-1} = d^2 q^{k-1}.$$

So (3) implies (1). □

Lemma C.1 shows that the condition of absolute irreducibility in the definition of (n, k, d) algebraic sources (Definition 1.2) is useful as it guarantees the existence of enough rational points. On the other hand, even if none of the irreducible components of V is absolutely irreducible, it may still be possible that the variety V has a substantial number of rational points, so that the randomness extraction question is meaningful.

Example. Let q be an odd prime power, and let a be a non-square in $\mathbb{F}_q$, i.e., $a \in \mathbb{F}_q^\times \setminus (\mathbb{F}_q^\times)^2$. Let $r \geq 2$. The affine variety $V = V(X_1^2 - aX_2^2) \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ over $\mathbb{F}_q$ is irreducible but not absolutely irreducible as $V_{\overline{\mathbb{F}}_q}$ consists of the two hyperplanes of $\mathbb{A}^r_{\overline{\mathbb{F}}_q}$ defined by $X_1 + \sqrt{a}X_2$ and $X_1 - \sqrt{a}X_2$ respectively. Let $W = V(X_1, X_2) \subseteq \mathbb{A}^r_{\mathbb{F}_q}$, which is an affine subspace of codimension two and is absolutely irreducible. Then $V(\mathbb{F}_q) = W(\mathbb{F}_q)$ and hence $|V(\mathbb{F}_q)| = q^{r-2}$.

There is a general algorithm that, given an affine variety $V \subseteq \mathbb{A}^r_{\mathbb{F}_q}$ over $\mathbb{F}_q$ such that none of the irreducible components of V of top dimension (i.e., dimension $\dim V$) is absolutely irreducible, outputs an affine subvariety W of lower dimension such that $W(\mathbb{F}_q) = V(\mathbb{F}_q)$ and W has an irreducible component of top dimension that is absolutely irreducible. Namely, for every irreducible component V_0 of V of top dimension, replace V_0 by the intersection of the irreducible components of $(V_0)_{\overline{\mathbb{F}}_q}$. This yields an affine variety over $\mathbb{F}_q$ as it is fixed by the Frobenius map over $\mathbb{F}_q$. Repeating this process, we would eventually obtain an affine subvariety $W \subseteq V$ such that $W(\mathbb{F}_q) = V(\mathbb{F}_q)$ and at least one irreducible component of W of top dimension is absolutely irreducible, as required by Definition 1.2. One can then choose the parameters k and d such that $D = f(U_{V(\mathbb{F}_q)}) = f(U_{W(\mathbb{F}_q)})$ is an (n, k, d) algebraic source over $\mathbb{F}_q$.

However, the problem is that we do not have a good bound on $\deg W$ or d. The general bound we know on $\deg W$ is doubly exponential in $\dim V$. It seems to be an interesting question to understand how large $\deg W$ can be given other parameters such as $\deg V$, $\dim V$, and the dimension r of the ambient space that contains V.

Proof of Lemma 4.12. We prove Lemma 4.12 now. First, we need the following lemma.

Lemma C.2. *Let $f : \mathbb{A}^n \to \mathbb{A}^m$ be a polynomial map over an algebraically closed field $\mathbb{F}$. Let $V \subseteq \mathbb{A}^n$ be an irreducible affine variety over $\mathbb{F}$, and let $W = \overline{f(V)} \subseteq \mathbb{A}^m$. Then there exists an affine subspace $H \subseteq \mathbb{A}^n$ of codimension $\dim V - \dim W$ such that $\dim(H \cap V) = \dim W$ and $\overline{f(H \cap V)} = W$.*

Proof. Consider a general point $x \in W$ and let $t = \dim f|_V^{-1}(x)$. Then $t = \dim V - \dim W$ by the fiber dimension theorem (Theorem 4.4). Then there exists an affine subspace $H \subseteq \mathbb{A}^n$ of codimension t such that $H \cap f|_V^{-1}(x) \neq \emptyset$, $\dim(H \cap f|_V^{-1}(x)) = \dim f|_V^{-1}(x) - t = 0$, and $\dim(H \cap V) = \dim V - t = \dim W$. (This can be shown by, e.g., taking H to be the intersection of t general hyperplanes containing a fixed point of $f|_V^{-1}(x)$ and then applying Lemma 4.1.) Then

$$\dim(f|_{H\cap V}^{-1}(x)) = \dim(H \cap f|_V^{-1}(x)) = 0 = \dim(H \cap V) - \dim W.$$

Let $W' = \overline{f(H \cap V)}$. By the fiber dimension theorem (applied to the morphism $f|_{H\cap V} : H \cap V \to W'$), we must have $\dim W' = \dim W$. As V is irreducible, so is W. It follows that $W' = W$. □

For convenience, we restate Lemma 4.12 below before presenting its proof.

Lemma 4.12. *Let $V \subseteq \mathbb{A}^n_{\mathbb{F}}$ be an affine variety over a field $\mathbb{F}$. Let $h_1, \ldots, h_s \in \mathbb{F}[X_1, \ldots, X_n]$ with $\deg h_1 \geq \cdots \geq \deg h_s$. Let $f_1, \ldots, f_m \in \mathcal{L}_{h_1,\ldots,h_s,\mathbb{F}}$, which define a polynomial map $f : \mathbb{A}^n_{\mathbb{F}} \to \mathbb{A}^m_{\mathbb{F}}$. Finally, let $W = \overline{f(V)} \subseteq \mathbb{A}^m_{\mathbb{F}}$. Then*

$$\deg W \leq \deg V \cdot \prod_{i=1}^{\dim W} \deg h_i.$$

Proof. By replacing V with $V_{\overline{\mathbb{F}}}$, we may assume that $\mathbb{F}$ is algebraically closed. By considering the irreducible components of V individually, we may assume that V is irreducible. Then W is also irreducible. Also, noting that performing an invertible linear transformation on $\mathbb{A}^m_{\mathbb{F}}$ (and replacing $f_1, \dots, f_m$ by their linear combinations accordingly) does not change the degrees of the subvarieties of $\mathbb{A}^m_{\mathbb{F}}$. So by Gaussian elimination, we may assume $f_i \in \mathcal{L}_{h_i,\dots,h_s,\mathbb{F}}$ for $i \in [m]$, where we let $\mathcal{L}_{h_i,\dots,h_s,\mathbb{F}} = \mathbb{F}$ if $i > s$.

Next, we reduce to the case where $\dim V = \dim W$. By Lemma C.2, there exists an affine subspace $H \subseteq \mathbb{A}^n_{\mathbb{F}}$ of codimension $\dim V - \dim W$ such that $\dim(H \cap V) = \dim W$ and $\overline{f(H \cap V)} = W$. Also note that $\deg(H \cap V) \le \deg V$ by Bézout's inequality (Lemma 4.2). Thus, by replacing V with $H \cap V$, we may assume $\dim V = \dim W$.

Let $k = \dim W$. By the fiber dimension theorem (Theorem 4.4), there exists a dense open subset U of W contained in $f(V)$ such that $\dim f|_V^{-1}(b) = 0$ for all $b \in U$. Let $B = W \setminus U$, which is a proper subvariety of W. As W is irreducible, we have $\dim B < k$. A general affine subspace $L \subseteq \mathbb{A}^m$ of codimension k then satisfies $L \cap B = \emptyset$ and $|L \cap W| = \deg W$. Fix such L. As $L \cap B = \emptyset$. We have $L \cap W \subseteq U \subseteq f(V)$. It follows that $L \cap W = f(f|_V^{-1}(L))$. Moreover, as $\dim f|_V^{-1}(b) = 0$ for all $b \in U$ and $L \cap W$ is a finite subset of U, the set $f|_V^{-1}(L) = f|_V^{-1}(L \cap W)$ is a finite set.

Suppose L is defined by degree-1 polynomials $\ell_1, \dots, \ell_k$. Then $f^{-1}(L)$ is defined by the polynomials $\ell_1(f_1, \dots, f_m), \dots, \ell_k(f_1, \dots, f_m)$. By Gaussian elimination, we may assume that for each $i \in [k]$, $\ell_i(f_1, \dots, f_m)$ does not involve the polynomials $f_1, \dots, f_{i-1}$, and hence $\ell_i(f_1, \dots, f_m) \in \mathcal{L}_{h_i,\dots,h_s,\mathbb{F}}$. In particular, we have $\deg \ell_i(f_1, \dots, f_m) \le \deg h_i$ for $i \in [k]$.

By Bézout's inequality, we have $|f|_V^{-1}(L)| = |f^{-1}(L) \cap V| \le \deg V \cdot \prod_{i=1}^k \deg h_i$. Therefore,

$$\deg(W) = |L \cap W| = |f(f|_V^{-1}(L))| \le |f|_V^{-1}(L)| \le \deg V \cdot \prod_{i=1}^k \deg h_i$$

as desired. □

Proof of Lemma 8.4. We now prove Lemma 8.4, which is restated below.

Lemma 8.4. *Let $C_0 \subseteq \mathbb{A}^n$ be an irreducible affine curve of degree d over an algebraically closed field $\mathbb{F}$. Then there exists an $\mathbb{F}$-linear field embedding $\tau : \mathbb{F}(C_0) \hookrightarrow \mathbb{F}((T))$ such that for any polynomial $f \in \mathbb{F}[X_1, \dots, X_n]$ of degree d that is not constant on C_0, the map τ sends f to $\tilde{f} \in \mathbb{F}((T))$ such that*

$$-\deg(C_0) \cdot d \le \operatorname{ord}(\tilde{f}) < 0.$$

Let C_0 be as in Lemma 8.4. Regard the affine space $\mathbb{A}^n$ as an open subset of the projective space $\mathbb{P}^n$ via $(a_1, \dots, a_n) \mapsto (1, a_1, \dots, a_n)$. Let C be the *projective closure* of $C_0 \subseteq \mathbb{A}^n$ in $\mathbb{P}^n$, which is an irreducible projective curve over $\mathbb{F}$ whose degree in $\mathbb{P}^n$ equals $\deg C_0$ (cf. Section A).

For a point $y \in C$, the field $\mathbb{F}(C)$ of rational functions on C has a subring $\mathcal{O}_{y,C}$, called the *local ring* of C at y, which consists of the rational functions that have no pole at y [LL89, Section 5]. It has a unique maximal ideal $\mathfrak{m}_{y,C}$, which consists of the rational functions on C that vanishes at y.

Let $\pi : \widetilde{C} \to C$ be the *normalization* of C, which is a finite and surjective morphism from a smooth and irreducible projective curve $\widetilde{C}$ to C [Sha94, §II.5].

We need the following facts.

Lemma C.3 ([LL89, Section 5]). *For $x \in \widetilde{C}$ and $y = \pi(x) \in C$, there exists a field embedding $\rho_x : \mathbb{F}(C) \hookrightarrow \mathbb{F}((T))$ such that $\rho_x(\mathcal{O}_{y,C}) \subseteq \mathbb{F}[[T]]$ and $\rho_x(\mathfrak{m}_{y,C}) \subseteq T \cdot \mathbb{F}[[T]]$.*

In addition, the following statement holds for the maps ρ_x chosen in [LL89, Section 5].

Lemma C.4 ([LL89, Proposition 3])**.** *Let $x \in \widetilde{C}$ and $y = \pi(x) \in C$. Let f and g be linear homogeneous polynomials nonzero on C such that $g(y) \neq 0$. Then f/g restricted to C is in $\mathcal{O}_{y,C}$, and $\operatorname{ord}(\rho_x(f/g)) \leq \deg C$.*

Proof of Lemma 8.4. Let $f \in \mathbb{F}[X_1, \ldots, X_n]$ be a polynomial of degree d. Identify X_i with Y_i/Y_0 for $i \in [n]$, where $Y_0, \ldots, Y_n$ are the n homogeneous coordinates of $\mathbb{P}^n$. Consider arbitrary $x \in \widetilde{C}$. We first show that

$$\operatorname{ord}(\rho_x(f)) \geq -\deg(C_0) \cdot d.$$

As f is a polynomial of degree d in $Y_1/Y_0, \ldots, Y_n/Y_0$, it suffices to show that $\operatorname{ord}(\rho_x(Y_i/Y_0)) \geq -\deg C_0$ for $i \in [n]$. Fix $i \in [n]$. Choose $j \in \{0, 1, \ldots, n\}$ such that Y_j does not vanish at y. Such an index j exists as $Y_0, \ldots, Y_n$ do not simultaneously vanish on $\mathbb{P}^n$. By Lemma C.3 and Lemma C.4, we have

$$0 \leq \operatorname{ord}(\rho_x(Y_i/Y_j)), \operatorname{ord}(\rho_x(Y_0/Y_j)) \leq \deg C = \deg C_0.$$

Then

$$\operatorname{ord}(\rho_x(Y_i/Y_0)) = \operatorname{ord}(\rho_x(Y_i/Y_j)) - \operatorname{ord}(\rho_x(Y_0/Y_j)) \geq -\deg C_0,$$

as desired.

Now assume that $f \in \mathbb{F}[X_1, \ldots, X_n]$ is not constant on C_0. It remains to show that there exists $x \in \widetilde{C}$ such that $\operatorname{ord}(\rho_x(f)) < 0$. This follows from the following standard argument: Assume to the contrary that $\operatorname{ord}(\rho_x(f)) \geq 0$ for all $x \in \widetilde{C}$. Then f is a regular function on $\widetilde{C}$ and hence defines a morphism from $\widetilde{C}$ to $\mathbb{A}^1$. Viewing $\mathbb{A}^1$ as an open subset of $\mathbb{P}^1$, we get a morphism $\varphi : \widetilde{C} \to \mathbb{P}^1$ whose image is contained in $\mathbb{A}^1$. On the other hand, it is well-known that the image of a morphism from a projective variety is closed [Sha94, §I.5.2, Theorem 2]. And as $\widetilde{C}$ is irreducible, we know $\varphi(\widetilde{C})$ is also irreducible. The only irreducible closed subsets of $\mathbb{P}^1$ are single points and $\mathbb{P}^1$ itself. As f is not constant on C_0, we know $\varphi(\widetilde{C})$ is not a single point. So $\varphi(\widetilde{C}) = \mathbb{P}^1$, which contradicts the fact $\varphi(\widetilde{C}) \subseteq \mathbb{A}^1$. □

Remark. The last part of the above proof can be strengthened to show that for a nonzero rational function g on C,

$$\sum_{x \in \widetilde{C} : \operatorname{ord}(\rho_x(g)) \neq 0} \operatorname{ord}(\rho_x(g)) = 0,$$

i.e., g has as many zeros as poles on $\widetilde{C}$, counting multiplicities. This implies the existence of $x \in \widetilde{C}$ such that $\operatorname{ord}(\rho_x(f)) < 0$ as follows: As f is not constant on C_0, there exists $c \in \mathbb{F}$ such that $g := f - c$ is not identically zero on C_0 but has at least one zero. Then g must also have a pole on $\widetilde{C}$, i.e., $\operatorname{ord}(\rho_x(g)) < 0$ for some $x \in \widetilde{C}$. Then as $f = g + c$, we have $\operatorname{ord}(\rho_x(f)) = \min\{\operatorname{ord}(\rho_x(g)), \operatorname{ord}(\rho_x(c))\} = \operatorname{ord}(\rho_x(g)) < 0$.

D Explicit Noether Normalization

We first prove Theorem 11.1 in the case where $\mathbb{F}$ is algebraically closed.

Proof of Theorem 11.1 when $\mathbb{F}$ is algebraically closed. Assume that $\mathbb{F}$ is algebraically closed. For each $u \in \mathbb{N}$, identify $\mathbb{A}^u$ with an open subset of $\mathbb{P}^u$ via $(x_1, \ldots, x_u) \mapsto (1, x_1, \ldots, x_u)$. Define

$$\Gamma = \{(x, y) \in V \times \mathbb{A}^m : y = \varphi(x)\} \subseteq \mathbb{A}^n \times \mathbb{A}^m \subseteq \mathbb{P}^n \times \mathbb{P}^m.$$

Let $\widetilde{\Gamma}$ be the (Zariski-)closure of Γ in $\mathbb{P}^n \times \mathbb{P}^m$. Then $\widetilde{\Gamma} \cap (\mathbb{A}^n \times \mathbb{A}^m) = \Gamma$ as Γ is closed in $\mathbb{A}^n \times \mathbb{A}^m$.

Let $\iota : V \to \Gamma$ be the morphism $x \mapsto (x, \varphi(x))$, which is an isomorphism between the affine varieties V and Γ over $\mathbb{F}$. Let $\pi_1 : \widetilde{\Gamma} \to \mathbb{P}^n$ and $\pi_2 : \widetilde{\Gamma} \to \mathbb{P}^m$ be the projections from $\widetilde{\Gamma}$ to the first factor and the second factor respectively. Then $\pi_2|_\Gamma \circ \iota = \varphi|_V$.

Claim D.1. *If $\pi_2^{-1}(\mathbb{A}^m) \subseteq \Gamma$, then $\varphi|_V$ is a finite morphism.*

To prove Claim D.1, we need a result from algebraic geometry. For a closed set Z of $\mathbb{A}^n \times \mathbb{A}^m$, the projection from Z to $\mathbb{A}^m$ is an example of an *affine morphism* to $\mathbb{A}^m$, making Z an *affine variety over* $\mathbb{A}^m$. Similarly, for a closed set Z of $\mathbb{P}^n \times \mathbb{A}^m$, the projection from Z to $\mathbb{A}^m$ is an example of a *projective morphism* to $\mathbb{A}^m$, making Z a *projective variety over* $\mathbb{A}^m$. See [Vak22] for the definitions of these objects with various degrees of generality. We need the following fact.

Lemma D.2. *A morphism to a variety is affine and projective iff it is finite.*

See [Vak22, Corollay 19.1.6]. The projectivity of π can be relaxed to *properness* [Gro61, Proposition 4.4.2]. To shed some light on Lemma D.2, consider an affine and projective morphism from a variety Z to $\mathbb{A}^0$ (i.e., a point) over $\mathbb{F}$, which just means that Z is an affine and projective variety over $\mathbb{F}$. In this case, the lemma simply states that the coordinate ring $\mathbb{F}[Z]$ is a finite-dimensional vector space over $\mathbb{F}$. As Z is affine, the set of regular functions (i.e., functions with no poles) on Z is $\mathbb{F}[Z]$. On the other hand, it is well-known that any regular function on a connected projective variety over $\mathbb{F}$ is constant. It follows that the dimension of $\mathbb{F}[Z]$ is indeed finite and equals the number of connected components of Z. (In fact, Z is just a finite collection of points.) Lemma D.2 states that finiteness holds more generally for any variety as the target of the morphism.

Proof of Claim D.1. The map $\pi_2|_\Gamma : \Gamma \to \mathbb{A}^m$ is an affine morphism. Suppose $\pi_2^{-1}(\mathbb{A}^m) \subseteq \Gamma$. Then $\Gamma = \widetilde{\Gamma} \cap (\mathbb{P}^n \times \mathbb{A}^m)$. So $\pi_2|_\Gamma$ is also a projective morphism. It follows from Lemma D.2 that $\pi_2|_\Gamma$ is finite. As ι is an isomorphism between V and Γ over $\mathbb{F}$, the map $\varphi|_V = \pi_2|_\Gamma \circ \iota$ is also finite. $\square$

By Claim D.1, we may assume that $\pi_2^{-1}(\mathbb{A}^m)$ contains a point $u \in \widetilde{\Gamma} \setminus \Gamma$. Let $x = \pi_1(u) \in \mathbb{P}^n$ and $y = \pi_2(u) \in \mathbb{A}^m$. Note that $x \notin \mathbb{A}^n$ as otherwise we would have $u = (x, y) \in \widetilde{\Gamma} \cap (\mathbb{A}^n \times \mathbb{A}^m) = \Gamma$. See the following diagrams for an illustration.

$$\begin{tikzcd}
\widetilde{\Gamma} \arrow[r, "\pi_1"] \arrow[d, "\pi_2"] & \mathbb{P}^n & u \arrow[r, mapsto, "\pi_1"] \arrow[d, mapsto, "\pi_2"] & x \in \mathbb{P}^n \setminus \mathbb{A}^n \\
\mathbb{P}^m & & y \in \mathbb{A}^m &
\end{tikzcd}$$

Claim D.3. *There exist $h_1(T), \ldots, h_n(T) \in \mathbb{F}((T))$ satisfying the following conditions:*

1. *At least one $h_i(T)$ has a pole, i.e., $h_i(T) \notin \mathbb{F}[[T]]$.*

2. *$P(h_1(T), \ldots, h_n(T)) = 0$ for all $P \in I(V)$.*

3. *$f_i(h_1(T), \ldots, h_n(T)) \in \mathbb{F}[[T]]$ for $i \in [m]$.*

Proof. We follow the argument in [KRS96]. As $\widetilde{\Gamma}$ is the closure of Γ in $\mathbb{P}^n \times \mathbb{P}^m$, the point $u \in \widetilde{\Gamma} \setminus \Gamma$ cannot possibly be an isolated point of $\widetilde{\Gamma}$ (i.e. a point that is an irreducible component of $\widetilde{\Gamma}$). So there exists an irreducible curve $C \subseteq \widetilde{\Gamma}$ that passes through u and intersects Γ. This can be shown by intersecting $\widetilde{\Gamma}$ with general hyperplanes containing u to reduce the dimension, and then picking an irreducible component of the intersection that contains u.

Let $\mathcal{O}_{u,C}$ be the local ring of C at u, and let $\mathfrak{m}_{u,C}$ be its unique maximal ideal. Then there exists a field embedding $\rho : \mathbb{F}(C) \hookrightarrow \mathbb{F}((T))$ such that $\rho(\mathcal{O}_{u,C}) \subseteq \mathbb{F}[[T]]$ and $\rho(\mathfrak{m}_{u,C}) \subseteq T \cdot \mathbb{F}[[T]]$, as we have seen in the proof of Lemma 8.4 in Appendix C.

Write $x = (x_0, \dots, x_n)$. As $x \in \mathbb{P}^n \setminus \mathbb{A}^n$. We have $x_0 = 0$ and $x_j \neq 0$ for some $j \in [n]$. Let $X_0, \dots, X_n$ be the homogeneous coordinates of $\mathbb{P}^n$. Then X_j does not vanish on x, and hence not on u either. It follows that X_i/X_j restricted to C is in $\mathcal{O}_{u,C}$ for $i = 0, \dots, n$. Let $g_i(T) = \rho(X_i/X_j) \in \mathbb{F}[[T]]$ for $i = 0, \dots, n$. As $\pi_1(\Gamma) \subseteq \mathbb{A}^n$, the function X_0 does not vanish on any point in Γ. As C intersects Γ, we see that X_0 does not vanish identically on C. So X_0/X_j restricts to a nonzero rational function on C and hence $g_0(T) = \rho(X_0/X_j) \neq 0$. For $i \in [n]$, let $h_i(T) = g_j(T)/g_0(T)$.

Recall that $g_0(T) = \rho(X_0/X_j)$. As $x_0 = 0$, we know X_0/X_j restricted to C is in $\mathfrak{m}_{u,C}$ and hence $g_0(T) \in T \cdot \mathbb{F}[[T]]$. And $g_j(T) = \rho(1) = 1$ by definition. So $h_j(T) = g_j(T)/g_0(T)$ has a pole, proving the first condition.

Consider arbitrary $P \in I(V)$. Let $\widetilde{P}(X_0, \dots, X_n) = X_0^{\deg(P)} P(X_1/X_0, \dots, X_n/X_0)$ be the homogenization of P. We know P vanishes identically on V and hence also on Γ. So $\widetilde{P}$ vanishes identically on its closure $\widetilde{\Gamma}$, which contains the curve C. So $\widetilde{P}(X_0/X_j, \dots, X_n/X_j)$ restricted to C is zero. It follows that $\widetilde{P}(g_0(T), \dots, g_n(T)) = \rho(\widetilde{P}(X_0/X_j, \dots, X_n/X_j)) = 0$. Therefore,

$$P(h_1(T), \dots, h_n(T)) = P(g_1(T)/g_0(T), \dots, g_n(T)/g_0(T)) = g_0(T)^{-\deg(P)} \widetilde{P}(g_0(T), \dots, g_n(T)) = 0,$$

proving the second condition.

Let $Y_1, \dots, Y_m$ be the coordinates of $\mathbb{A}^m$, which (restricted to C) are in $\mathcal{O}_{u,C}$ as $u \in \mathbb{P}^n \times \mathbb{A}^m$. Consider arbitrary $i \in [m]$. Let $F_i(X_1, \dots, X_n, Y_i) = f_i(X_1, \dots, X_n) - Y_i$. Let $\widetilde{F}_i(X_0, \dots, X_n, Y_i) = X_0^{\deg(f_i)}(f_i(X_1/X_0, \dots, X_n/X_0) - Y_i)$ be the homogenization of F_i with respect to $X_1, \dots, X_n$. By definition, F_i vanishes identically on $\Gamma \subseteq \mathbb{A}^n \times \mathbb{A}^m$ and hence $\widetilde{F}_i$ vanishes identically on $\widetilde{\Gamma} \cap (\mathbb{P}^n \times \mathbb{A}^m)$. As $C \cap (\mathbb{P}^n \times \mathbb{A}^m)$ is a subset of $\widetilde{\Gamma} \cap (\mathbb{P}^n \times \mathbb{A}^m)$ and is dense in C, we see that $\widetilde{F}(X_0/X_j, \dots, X_n/X_j, Y_i)$ restricted to C is zero. Then $\widetilde{F}_i(g_0(T), \dots, g_n(T), \rho(Y_i)) = \rho(\widetilde{F}(X_0/X_j, \dots, X_n/X_j, Y_i)) = 0$. Therefore,

$$\begin{aligned} f_i(h_1(T), \dots, h_n(T)) - \rho(Y_i) &= F_i(h_1(T), \dots, h_n(T), \rho(Y_i)) \\ &= F_i(g_1(T)/g_0(T), \dots, g_n(T)/g_0(T), \rho(Y_i)) \\ &= g_0(T)^{-\deg(f_i)} \widetilde{F}_i(g_0(T), \dots, g_n(T), \rho(Y_i)) \\ &= 0. \end{aligned}$$

So $f_i(h_1(T), \dots, h_n(T)) = \rho(Y_i) \in \rho(\mathcal{O}_{u,C}) \subseteq \mathbb{F}[[T]]$, proving the third condition. □

However, Lemma 6.8 states that Claim D.3 cannot be true. So we obtain a contradiction, implying that the assumption that $\pi_2^{-1}(\mathbb{A}^m)$ contains a point $u \in \widetilde{\Gamma} \setminus \Gamma$ is false. This concludes the proof of Theorem 11.1 when $\mathbb{F}$ is algebraically closed.

Proof of Theorem 11.1 for arbitrary $\mathbb{F}$. Theorem 11.1 for an aritrary field $\mathbb{F}$ follows from the special case where $\mathbb{F}$ is algebraically closed and the fact that finiteness descends under a "faithfully flat base change." We explain this now.

In the following, all rings and algebras are commutative with unity.

Let R be a ring. For R-modules M and M', their *tensor product* $M \otimes_R M'$ over R is defined to be the R-module generated by the set of elements $\{a \otimes b : a \in M, b \in M'\}$ subject to the R-bilinear relations $a \otimes b + a' \otimes b = (a + a') \otimes b$, $a \otimes b + a \otimes b' = a \otimes (b + b')$, and $c(a \otimes b) = (ca) \otimes b = a \otimes (cb)$ for $a, a' \in M$, $b, b' \in M'$, and $c \in R$. For an R-algebra S and an R-module M, the tensor product $M \otimes_R S$ is an S-module.

An R-module is *free* if it has a generating set that is linearly independent over R. We need the following fact about nonzero free modules (which holds more generally for *faithfully flat modules* [Vak22, Definition 25.5.1]).

Lemma D.4. *Let F be a nonzero free R-module. Let $M \to M'$ be an R-module homomorphism. Then $M \to M'$ is surjective iff the induced map $M \otimes_R F \to M' \otimes_R F$ is surjective.*

Lemma D.4 implies the following fact about descent of finiteness.

Lemma D.5. *Let S be an R-algebra that is also a nonzero free R-module. Let M be an R-module. Suppose $M \otimes_R S$ is a finitely generated S-module. Then M is a finitely generated R-module.*

Proof. Let $\{r_1, \ldots, r_n\}$ be a finite set of generators of $M \otimes_R S$. By definition, we may write each r_i as a finite sum $r_i = \sum_{j \in I_i} m_{ij} \otimes s_{ij}$ over an index set I_i with $m_{ij} \in M$ and $s_{ij} \in S$. Form the free R-module F with the basis $\{e_{ij} : i \in [n], j \in I_i\}$. Consider the R-module homomorphism $F \to M$ sending e_{ij} to m_{ij}. The induced map $F \otimes_R S \to M \otimes_R S$ is surjective as $\sum_{j \in I_i} e_{ij} \otimes s_{ij}$ is sent to r_i for $i \in [n]$. By Lemma D.4, the map $F \to M$ is also surjective. So M is generated by the finite set $\{m_{ij} : i \in [n], j \in I_i\}$ as an R-module. □

Now we are ready to prove Theorem 11.1 in full generality.

Proof of Theorem 11.1. Let $\widetilde{\varphi} = \varphi|_V : V \to \mathbb{A}^m_{\mathbb{F}}$, which is associated with the $\mathbb{F}$-algebra homomorphism $\widetilde{\varphi}^\sharp : \mathbb{F}[Y_1, \ldots, Y_m] \to \mathbb{F}[V]$. Let $R = \widetilde{\varphi}^\sharp(\mathbb{F}[Y_1, \ldots, Y_m]) \subseteq \mathbb{F}[V]$. By definition, we want to show that $\mathbb{F}[V]$ is a finitely generated module over R.

For an affine variety W over $\mathbb{F}$, the coordinate ring of $W_{\overline{\mathbb{F}}}$ may be identified with $\mathbb{F}[W] \otimes_{\mathbb{F}} \overline{\mathbb{F}}$. As we already know that Theorem 11.1 holds over algebraically closed fields by the discussion above, applying Theorem 11.1 to the morphism $V_{\overline{\mathbb{F}}} \to \mathbb{A}^m_{\overline{\mathbb{F}}}$ defined by $f_1, \ldots, f_m$ shows that $\mathbb{F}[V] \otimes_{\mathbb{F}} \overline{\mathbb{F}}$ is a finitely generated module over $R \otimes_{\mathbb{F}} \overline{\mathbb{F}}$.

The ring $R \otimes_{\mathbb{F}} \overline{\mathbb{F}}$ is a free module over R since $\overline{\mathbb{F}}$ is free over $\mathbb{F}$ and tensor products commute with direct sums. By Lemma D.5, to show that $\mathbb{F}[V]$ is a finitely generated module over R, it suffices to show that $\mathbb{F}[V] \otimes_R (R \otimes_{\mathbb{F}} \overline{\mathbb{F}})$ is a finitely generated module over $R \otimes_{\mathbb{F}} \overline{\mathbb{F}}$.

Finally, we have the canonical isomorphisms

$$\mathbb{F}[V] \otimes_R (R \otimes_{\mathbb{F}} \overline{\mathbb{F}}) \cong (\mathbb{F}[V] \otimes_R R) \otimes_{\mathbb{F}} \overline{\mathbb{F}} \cong \mathbb{F}[V] \otimes_{\mathbb{F}} \overline{\mathbb{F}}.$$

See [AM69, Proposition 2.14]. So $\mathbb{F}[V] \otimes_R (R \otimes_{\mathbb{F}} \overline{\mathbb{F}})$ is a finitely generated module over $R \otimes_{\mathbb{F}} \overline{\mathbb{F}}$, as desired. □

Proof of Theorem 11.2. Theorem 11.2 is almost equivalent to Theorem 11.1 except that the coordinate ring of an affine variety is always reduced (i.e., it has no nonzero nilpotent elements) while the algebra A in Theorem 11.2 may be non-reduced. Nevertheless, we can easily derive Theorem 11.2 from Theorem 11.1 as follows.

Let R' be an algebra over a ring R. We say R' is *integral* over R if every $b \in R'$ is a root of a *monic* polynomial $X^t + a_{t-1}X^{t-1} + \cdots + a_0$ with the coefficients $a_i \in R$. We need the following fact, which states that integrality is equivalent to finiteness for finitely generated algebras.

Lemma D.6 ([AM69, Proposition 5.1 and Corollary 5.2])**.** *A finitely generated algebra over a ring R is a finitely generated module over R iff it is integral over R.*

Now we give the proof of Theorem 11.2.

Proof of Theorem 11.2. If I is radical, then $I = I(V(I))$ and hence $A \cong \mathbb{F}[V(I)]$. In this case, the theorem follows from Theorem 11.1 and the definition of finite morphisms.

Now consider general I. Identify A with $\mathbb{F}[X_1, \ldots, X_n]/I$ and let $\bar{A} = \mathbb{F}[X_1, \ldots, X_n]/\mathrm{rad}(I) = A/\mathrm{nil}(A)$, where $\mathrm{rad}(I)$ denotes the radical of I and $\mathrm{nil}(A)$ denotes the nilradical of A, i.e., the radical of the zero ideal. As $\mathrm{rad}(I)$ is radical and $V(\mathrm{rad}(I)) = V(I)$, we see that $\bar{A}$ is a finitely generated module over S by Theorem 11.1. So $\bar{A}$ is integral over S by Lemma D.6. Now consider arbitrary $b \in A$. As $\bar{A}$ is integral over S, there exists a monic polynomial $F \in S[X]$ such that the image of $F(b)$ in $\bar{A}$ is zero, or equivalently, $F(b) \in \mathrm{nil}(A)$. So $F(b)^N = 0$ for some $N \in \mathbb{N}^+$. Then b is a root of the monic polynomial F^N. Therefore, A is integral over S and hence is a finitely generated module over S by Lemma D.6. □

References

[ABGS21] Divesh Aggarwal, Huck Bennett, Alexander Golovnev, and Noah Stephens-Davidowitz. Fine-grained hardness of CVP(P) - everything that we can prove (and nothing else). In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021*, pages 1816–1835. SIAM, 2021.

[AM69] Michael F. Atiyah and I. G. MacDonald. *Introduction to Commutative Algebra.* Addison-Wesley-Longman, 1969.

[BCS97] Peter Bürgisser, Michael Clausen, and Mohammad A. Shokrollahi. *Algebraic Complexity Theory.* Springer, 1997.

[BDL16] Jean Bourgain, Zeev Dvir, and Ethan Leeman. Affine extractors over large fields with exponential error. *computational complexity*, 25(4):921–931, 2016.

[BE19] Juliette Bruce and Daniel Erman. A probabilistic approach to systems of parameters and noether normalization. *Algebra & Number Theory*, 13(9):2081–2102, 2019.

[Bom66] Enrico Bombieri. On exponential sums in finite fields. *Amer. J. Math.*, 88:71–105, 1966.

[Bou07] Jean Bourgain. On the construction of affine extractors. *GAFA Geometric And Functional Analysis*, 17(1):33–57, 2007.

[CGL21] Eshan Chattopadhyay, Jesse Goodman, and Jyun-Jie Liao. Affine extractors for almost logarithmic entropy. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2021.

[CM06] Antonio Cafure and Guillermo Matera. Improved explicit estimates on the number of solutions of equations over a finite field. *Finite Fields and Their Applications*, 12(2):155–185, 2006.

[CZ19] Eshan Chattopadhyay and David Zuckerman. Explicit two-source extractors and resilient functions. *Annals of Mathematics*, 189:653–705, 2019.

[Del74] Pierre Deligne. La conjecture de Weil. I. *Publications Mathématiques de l'Institut des Hautes Études Scientifiques*, 43(1):273–307, 1974.

[Del78] Philippe Delsarte. Bilinear forms over a finite field, with applications to coding theory. *Journal of combinatorial theory, Series A*, 25(3):226–241, 1978.

[DGW09] Zeev Dvir, Ariel Gabizon, and Avi Wigderson. Extractors and rank extractors for polynomial sources. *computational complexity*, 18(1):1–58, 2009.

[DKL14] Zeev Dvir, János Kollár, and Shachar Lovett. Variety evasive sets. *computational complexity*, 23(4):509–529, 2014.

[DL78] Richard A. Demillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7(4):193–195, 1978.

[Dvi12] Zeev Dvir. Extractors for varieties. *computational complexity*, 21(4):515–572, 2012.

[FG15] Michael A. Forbes and Venkatesan Guruswami. Dimension expanders via rank condensers. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015*, volume 40 of *LIPIcs*, pages 800–814. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.

[For14] Michael A. Forbes. *Polynomial identity testing of read-once oblivious algebraic branching programs.* PhD thesis, Massachusetts Institute of Technology, 2014.

[FS12] Michael A. Forbes and Amir Shpilka. On identity testing of tensors, low-rank recovery and compressed sensing. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, pages 163–172, 2012.

[Gab85] Ernest M. Gabidulin. Theory of codes with maximum rank distance. *Problemy Peredachi Informatsii*, 21(1):3–16, 1985.

[GKW21] Alexander Golovnev, Alexander S. Kulikov, and R. Ryan Williams. Circuit depth reductions. In *12th Innovations in Theoretical Computer Science Conference, ITCS 2021*, volume 185 of *LIPIcs*, pages 24:1–24:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[GR08] Ariel Gabizon and Ran Raz. Deterministic extractors for affine sources over large fields. *Combinatorica*, 28(4):415–440, 2008.

[Gro61] Alexander Grothendieck. Éléments de géométrie algébrique: III. Étude cohomologique des faisceaux cohérents, Première partie. *Publications Mathématiques de l'IHÉS*, 11:5–167, 1961.

[GRS06] Ariel Gabizon, Ran Raz, and Ronen Shaltiel. Deterministic extractors for bit-fixing sources by obtaining an independent seed. *SIAM Journal on Computing*, 36(4):1072–1094, 2006.

[GSS19] Zeyu Guo, Nitin Saxena, and Amit Sinhababu. Algebraic dependencies and PSPACE algorithms in approximative complexity over any field. *Theory of Computing*, 15(1):1–30, 2019.

[GW97] Oded Goldreich and Avi Wigderson. Tiny families of functions with random properties: A quality-size trade-off for hashing. *Random Structures & Algorithms*, 11(4):315–343, 1997.

[Hei83] Joos Heintz. Definability and fast quantifier elimination in algebraically closed fields. *Theoretical Computer Science*, 24(3):239–277, 1983.

[HR15] Pavel Hrubeš and Anup Rao. Circuits with medium fan-in. In *30th Conference on Computational Complexity, CCC 2015, June 17-19, 2015, Portland, Oregon, USA*, volume 33 of *LIPIcs*, pages 381–391. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.

[HS80] Joos Heintz and Claus-Peter Schnorr. Testing polynomials which are easy to compute. In *Proceedings of the twelfth annual ACM Symposium on Theory of Computing*, pages 262–272, 1980.

[Jel05] Zbigniew Jelonek. On the effective Nullstellensatz. *Inventiones mathematicae*, 162(1):1–17, 2005.

[KPS01] Teresa Krick, Luis Miguel Pardo, and Martín Sombra. Sharp estimates for the arithmetic Nullstellensatz. *Duke Mathematical Journal*, 109(3):521–598, 2001.

[KRS96] János Kollár, Lajos Rónyai, and Tibor Szabó. Norm-graphs and bipartite Turán numbers. *Combinatorica*, 16(3):399–406, 1996.

[Lan02] Serge Lang. *Algebra*. Springer-Verlag, New York Inc., third edition, 2002.

[Li11] Xin Li. A new approach to affine extractors and dispersers. In *Proceedings of the 26th Annual IEEE Conference on Computational Complexity, CCC 2011, San Jose, California, USA, June 8-10, 2011*, pages 137–147. IEEE Computer Society, 2011.

[Li16] Xin Li. Improved two-source extractors, and affine extractors for polylogarithmic entropy. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016*, pages 168–177. IEEE Computer Society, 2016.

[LL89] Thomas Lehmkuhl and Thomas Lickteig. On the order of approximation in approximative triadic decompositions of tensors. *Theoretical Computer Science*, 66(1):1–14, 1989.

[LZ19] Fu Li and David Zuckerman. Improved extractors for recognizable and algebraic sources. In *23rd International Conference on Randomization and Computation (RANDOM)*, 2019.

[Mat22] Guillermo Matera. Personal communication, 2022.

[MK93] Oscar Moreno and P Vijay Kumar. Minimum distance bounds for cyclic codes and Deligne's theorem. *IEEE Transactions on Information Theory*, 39(5):1524–1534, 1993.

[MS77] Florence Jessie MacWilliams and Neil James Alexander Sloane. *The Theory of Error Correcting Codes*, volume 16. Elsevier, 1977.

[Nag53] Masayoshi Nagata. Some remarks on local rings. *Nagoya Mathematical Journal*, 6:53–58, 1953.

[Nag56] Masayoshi Nagata. A general theory of algebraic geometry over Dedekind domains, I: the notion of models. *American Journal of Mathematics*, 78(1):78–116, 1956.

[Nag62] Masayoshi Nagata. *Local Rings*. New York, Interscience Publishers, 1962.

[NN93] Joseph Naor and Moni Naor. Small-bias probability spaces: efficient constructions and applications. *SIAM J. Comput.*, 22(4):838–856, 1993.

[Noe26] Emmy Noether. Der endlichkeitssatz der invarianten endlicher linearer gruppen der charakteristik p. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1926:28–35, 1926.

[Per27] Oskar Perron. *Algebra I (Die Grundlagen)*. W. de Gruyter, Berlin, 1927.

[Rao07] Anup Rao. An exposition of Bourgain's 2-source extractor. In *TR 07-034*. Electronic Colloqium on Computaitonal Complexity, 2007.

[Rem16] Zachary Remscrim. The Hilbert function, algebraic extractors, and recursive Fourier sampling. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 197–208. IEEE, 2016.

[Rob83] Guy Robin. Estimation de la fonction de Tchebychef θ sur le k-ième nombre premier et grandes valeurs de la fonction ω (n) nombre de diviseurs premiers de n. *Acta Arithmetica*, 42:367–389, 1983.

[Sch80] Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4):701–717, 1980.

[Sha94] Igor R. Shafarevich. *Basic Algebraic Geometry 1: Varieties in Projective Space*. Springer-Verlag, 1994.

[Sha08] Ronen Shaltiel. How to get more mileage from randomness extractors. *Random Structures & Algorithms*, 33(2):157–186, 2008.

[Sho09] Victor Shoup. *A computational introduction to number theory and algebra*. Cambridge university press, 2009.

[Vad12] Salil Vadhan. Pseudorandomness. *Foundations and Trends® in Theoretical Computer Science*, 7(1-3):1–336, 2012.

[Vak22] Ravi Vakil. The Rising Sea: Foundations of Algebraic Geometry. `https://math.stanford.edu/~vakil/216blog/FOAGaug2922publici.pdf`, 2022. August 29, 2022 version.

[Yeh11] Amir Yehudayoff. Affine extractors over prime fields. *Combinatorica*, 31(2):245–256, 2011.

[Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. *Symbolic and Algebraic Computation*, pages 216–226, 1979.